Edge-connectivity of undirected and directed hypergraphs

by

Tamás Király

Dissertation submitted to Eötvös Loránd University,
Faculty of Science, for the Ph. D. degree

Doctoral School: Mathematics
Director: Miklós Laczkovich

Doctoral Program: Applied Mathematics
Director: András Prékopa

Supervisor: András Frank
Professor, Doctor of Sciences

The dissertation was prepared at the Operations Research Deptartment of Eötvös Loránd University and at the Egerváry Research Group on Combinatorial Optimization (MTA-ELTE)

Budapest, May 2003

Contents

Acknowledgements

I am greatly indebted to my supervisor, András Frank, who introduced me to this area of combinatorial optimization, and who helped me all the time by providing guidance and support, and by sharing his experience and unique understanding of the subject. His ideas and suggestions were essential for every new result of the thesis, and it was (and is) a real privilege for me to work with him.

I am grateful to Zoltán Király and Márton Makai, with whom I worked on many of the results featured in the thesis. Apart from contributing their own results, they also helped a lot by suggesting changes to some of my less comprehensible proofs. For the same reason I am grateful to the unknown referees of our papers.

I would like to thank every member of the Egerváry Research Group for the excellent atmosphere and highly motivating discussions. It is a rare joy to work in such a stimulating and enjoyable environment.

Thanks are also due to Matthias Kriesell for his contribution, and to Katherina Langkau, Gábor Maróti, Satoru Iwata, and Alex Berg for fruitful discussions on some topics included in the thesis. Let me also express my thanks the staff of the Operations Research Department of Eötvös Loránd University, who provided invaluable help on many occasions.

Finally, I am grateful to my family for giving every support I needed, and for tolerating my unusual working habits.

Chapter 1

Introduction and preliminaries

1.1 Overview

1.1.1 Objectives

The area of connectivity problems is one of the areas of graph theory that have both considerable theoretical interest and an important range of practical applications. The mathematical notion of connectivity appears naturally for example in the design of telecommunication networks and in VLSI design. As these technologies develop at a fast rate and require the solution of more and more complex optimization problems, there is a growing interest in the corresponding mathematical concepts and algorithms. At the same time, these mathematical concepts are interesting in their own right, leading to deep and beautiful structural results. Graph connectivity, and in particular graph edge-connectivity is therefore a well-studied area of both theoretical and applied mathematics.

Hypergraphs are a natural generalization of graphs, and many models in telecommunication and networking make use of this more general notion. In this light the extension of connectivity problems to hypergraphs is a plausible step. This area is less intensively studied than graph connectivity, recent result show however that some deep and elegant generalizations of graph results can be obtained, sometimes by novel mathematical methods.

In this thesis we discuss problems that are related to the edge-connectivity of hypergraphs. An intuitive description of edge-connectivity problems is that they concern the robustness of networks with respect to the failure of (hyper)edges, in contrast to vertex-connectivity problems, where the failure of nodes is also considered.

The aim of the thesis is twofold: on one hand, to extend the techniques used in graph edge-connectivity (for example the splitting-off technique, or exploitation of relations between undirected and directed graphs) to hypergraphs, and on the other hand, to present a framework for combining connectivity augmentation and orientation problems. Both of these themes are addressed with a strong emphasis on the role of submodularity and related properties intrinsic to the structure of these problems.

The understanding of the links to submodularity provides an important tool for obtaining structural descriptions and min-max theorems for the different variants of edge-connectivity properties for graphs, hypergraphs and directed hypergraphs. Moreover, the well-known algorithmic properties of submodular systems imply that efficient algorithms can be constructed based on the structural results. The precise analysis and optimization of these algorithms is beyond the scope of the thesis, usually only the existence of polynomial time algorithms will be stated.

There is a lot of ongoing research in the field, most of which is not addressed here in detail (for example the very interesting class of problems featuring both connectivity and parity), so the thesis is by no means a comprehensive survey. Rather, it intends to provide a concise account of connections and structural relations concerning a well-defined class of problems. This includes several new results on hypergraph edge-connectivity, some of which are interesting even when specialized to graphs.

1.1.2 Edge-connectivity – brief history

The most basic properties of edge-connectivity were established by different versions of Menger's classical theorem on the relation between cuts and edge-disjoint paths, and by the concept of network flows [21]. Important structural results were obtained by Gomory and Hu [42] and Edmonds [16]. Another significant contribution to the development of the field was the study of packings and coverings by trees, where the fundamental results are due to Tutte [70] and Nash-Williams [63]. Analogous theorems for directed graphs were established by Edmonds [14].

These areas are closely linked to the theory of matroids and submodular functions. We do not give a detailed account of the many fundamental results on submodularity by Edmonds, Lovász and others; for a concise bibliography see [26]. Efficient frameworks for dealing with connectivity problems in the context of submodularity were given for example in [28] and [66]. One notion of particular importance is that of submodular flows, introduced by Edmonds and Giles [17], which is the most popular general algorithmic framework featuring submodular functions.

One major development of the research going on in the field was the successful work on edge-connectivity augmentation problems. Initial results were obtained by Eswaran and Tarjan [18], and Plesnik [64]. Fundamental deep results were due to Lovász [55] and Mader ([58], [59]) on the splitting-off operation and to Watanabe and Nakamura [71] who solved the problem of augmenting a graph with minimal number of edges to make it k-edge-connected. Frank [29] showed how strongly polynomial algorithms can be designed with the help of the splitting-off operation. Recently, the research on edge-connectivity of hypergraphs has also started to gain ground. We give a detailed account on the developments in this direction in Chapter 3.

Another branch of edge-connectivity problems concerns the orientation of graphs. It was Robbins [65] who first considered the problem of finding orientations of graphs with certain connectivity properties. One of the fundamental results of the area is Nash-Williams' beautiful result on well-balanced orientations of graphs [61]. Frank [22] obtained important results on orientation problems that can be formulated using supermodular functions, and he showed how submodular flows can be used for this purpose [25].

An essential concern in relation to these problems is the construction of efficient algorithms. A lot of the above results yield polynomial algorithms, but often they are too cumbersome for practical purposes. There have been a lot of work on finding more efficient algorithms, for example by Gabow [39] and Benczúr [6].

The algorithmic aspect is particularly important because of the wide range of real-world applications. In telecommunications and informatics, guaranteeing the stability and robustness of networks leads to problems that can be modeled by some variation of graph or hypergraph connectivity. Of course, these problems are usually much more complex than the relatively simple formulations for which elegant mathematical structural characterizations exist, and in most cases they are NP-hard; but the algorithms and characterizations for the simple models can often be used as building blocks in approximate solutions for the real-world problems.

1.1.3 Structure of the thesis

The rest of this chapter contains preliminaries on notations and some fundamental results. First we present the different variants of edge-connectivity for graphs and digraphs. This is followed by a brief exposition of connectivity for hypergraphs, and the introduction of directed hypergraphs as a generalization of directed graphs. Finally, some definitions are given concerning families of sets, which will be used throughout the thesis in the proofs and in the formulation of min-max results.

In Chapter 2 we review results on submodular and supermodular functions that will be basic tools later in the thesis. Most characteristics are presented from a polyhedral point of view, and we discuss possible relaxations of submodularity that preserve these polyhedral characteristics. The section on relaxations contains some new results in this respect.

Chapter 3 is devoted to edge-connectivity augmentation problems. After a brief account of the state of the art for graphs, we present how augmentation problems can be extended to hypergraphs in different ways depending on the objective to be attained. These extensions differ in difficulty, and we cite an NP-completeness result to indicate that some problems are considerably more difficult than their graph counterpart. The main result of the chapter (based on [51]) concerns the construction of uniform hypergraphs covering given symmetric crossing supermodular set functions. This problem includes as a special case the k-edge-connectivity augmentation of hypergraphs by uniform hyperedges.

As a generalization of digraph augmentation problems, Chapter 4 discusses connectivity augmentation of directed hypergraphs, or in more general terms, covering crossing supermodular functions by directed hypergraphs (based on [52]). As in the undirected case, different objectives can be specified, and extensions of digraph augmentation can be formulated for many problems, including (k,l)-edge-connectivity augmentation. The main tool used here is a slight generalization of a splitting-off result of Berg, Jackson, and Jordán [7], which enables us to extend their results on k-edge-connectivity augmentation to more general connectivity requirements..

As it was already mentioned, submodularity and matroid theory plays an important role in the description of edge-connectivity problems. The most obvious link is between partition-connectivity and the circuit matroid of graphs. In Chapter 5, which presents the results in [37], we analyze some properties of the so-called hypergraphic matroid, introduced by Lorea [54], which is a direct generalization of the circuit matroid of graphs. This leads to a new type of connectivity notion for hypergraphs, generalizing (k,l)-partition-connectivity of graphs, and the matroid-theoretic approach gives an insight into its structural properties. At the end of the chapter it is shown how these results can be applied to the problem of finding disjoint Steiner trees of graphs.

After discussing connectivity properties of directed and undirected hypergraphs, it is natural to consider the possible extension of graph orientation problems related to connectivity. Chapter 6 starts with an account of graph orientation results, and the link between partition-connectivity of graphs and edge-connectivity of digraphs. Then we examine how the basic orientation results can be extended to hypergraphs, and we present some results on local requirements which are new even when they are specialized to graphs. Using these techniques, a link is established between (k,l)-partition-connectivity of hypergraphs and (k,l)-edge-connectivity of directed hypergraphs. The chapter also includes a new proof and an extension of a very interesting result of Khanna, Naor, and Shepherd [50] on network design with orientation constraints. The new results of the chapter appeared in [36].

Using the links established between undirected and directed connectivity problems, Chapter 7 examines a possibility for combining connectivity augmentation and orientation problems. In view of characterizations for graphs and hypergraphs that have an orientation covering certain set functions, it can be asked how many (hyper)edges must be added to an initial (hyper)graph to have such an orientation. This problem includes as a special case the (k, l)-partition-connectivity augmentation of graphs and hypergraphs. The theorems presented here are from [35] and [52]. In addition, some results are extended to mixed hypergraphs, where the characterizations are much more complicated.

1.2 Edge-connectivity and orientation of graphs

In the following pages we introduce the basic notions and notations that will be used throughout the thesis. We assume that the reader is familiar with the basics of graph theory, but for sake of clarity the fundamentals of graph connectivity and orientation are presented in some detail.

1.2.1 Undirected and directed graphs

We use the notation $\mathbb{Z}_+$ ($\mathbb{Q}_+$, $\mathbb{R}_+$) for the set of non-negative integer (rational, real) numbers. For $x \in \mathbb{R}$, $(x)^+$ denotes $\max\{x, 0\}$.

Let V be a finite ground set. For subsets X and Y of V, we use the notation $X - Y := \{v \in X : v \notin Y\}$. For a node $y \in V$, we write $X - y := X - \{y\}$ and $X + y := X \cup \{y\}$. The characteristic function of a set X will be denoted by $\chi_X : V \to \{0; 1\}$, i.e. $\chi_X(v) = 1$ if $v \in X$ and $\chi_X(v) = 0$ otherwise.

A graph $G = (V, E)$ or a directed graph $D = (V, A)$ (digraph for short) is allowed to have loops and parallel edges unless otherwise stated. For an edge a of a digraph we write $a = uv$ if u is the tail-node of a and v is the head-node.

The degree of a node $v \in V$ in G is denoted by $d_G(v)$, while $\varrho_D(v)$ denotes the in-degree of v in D and $\delta_D(v)$ is the out-degree. An edge in G or D is *induced* by a set $X \subseteq V$ if both of its endnodes are in X. An edge e of G enters X if exactly one of its endnodes is in X, while an edge $a \in A$ enters X if its head-node is in X but its tail-node is not. For a set $X \subseteq V$, we define $\Delta_G(X) = \{e \in E : e \text{ enters } X\}$, $\Delta^-_D(X) = \{a \in A : a \text{ enters } X\}$, and $\Delta^+_D(X) = \{a \in A : a \text{ enters } V - X\}$. The notations for the cardinalities of these edge-sets are: $d_G(X) := |\Delta_G(X)|$, $\varrho_D(X) := |\Delta^-_D(X)|$, and $\delta_D(X) = |\Delta^+_D(X)|$. The number of edges of G (or D) induced by a set $X \subseteq V$ is denoted by $i_G(X)$ (or $i_D(X)$). Sometimes we will indicate the edge set in the subscript, i.e. $i_E(X)$, $d_E(X)$, $\varrho_A(X)$, etc.

The *underlying graph* of a directed graph $D = (V, A)$ is obtained by replacing each directed edge $uv \in A$ by the undirected edge uv. An *orientation* of a graph $G = (V, E)$ is a directed graph $D = (V, A)$ whose underlying graph is G.

A *set function* is a function $p : 2^V \to \mathbb{Q} \cup \{-\infty, +\infty\}$ (sometimes the set function is defined only on a given family of subsets of V). It is assumed that $p(\emptyset) = 0$, unless otherwise stated. For a function $m : V \to \mathbb{Q}$ and a set $X \subseteq V$, we use the notation

$$m(X) := \sum_{v \in X} m(v).$$

A graph G is said to *cover* a set function p if $d_G(X) \geq p(X)$ for every $X \subseteq V$. Analogously, a digraph D covers p if $\varrho_D(X) \geq p(X)$ for every $X \subseteq V$. Given a function $m : V \to \mathbb{Z}_+$, the graph G satisfies the *degree-specification* m if $d_G(v) = m(v)$ for every $v \in V$. If we have an in-degree specification $m_i : V \to \mathbb{Z}_+$ and an out-degree specification $m_o : V \to \mathbb{Z}_+$, the digraph D satisfies the degree-specifications if $\varrho_D(v) = m_i(v)$ and $\delta_D(v) = m_o(v)$ for every $v \in V$.

In optimization problems we often consider a cost function $c : E \to \mathbb{Q}$ on the edges of the graph $(G = V, E)$ or $c : A \to \mathbb{Q}$ on the edges of a digraph $D = (V, A)$. A cost function on E is said to be *node induced* if $c(uv) = c'(u) + c'(v)$ where $c' : V \to \mathbb{Q}$ is a cost function on the nodes. Similarly, a cost function c on A is node induced if $c(uv) = c'_o(u) + c'_i(v)$ where $c'_i : V \to \mathbb{Q}$ and $c'_o : V \to \mathbb{Q}$ are cost functions on the nodes.

1.2.2 The notion of edge-connectivity

For two nodes $s \in V$ and $t \in V$, a set $X \subseteq V$ is called an *$\overline{s}t$-set* if $s \notin X$ and $t \in X$. Description of undirected and directed edge-connectivity is based on the following versions of Menger's Theorem (see e.g. [21]):

Theorem 1.1 (Menger). *Let $G = (V, E)$ be a graph, and $s, t \in V$ distinct nodes. The maximum number of edge-disjoint paths between s and t is*

$$\min\{d_G(X) : \ X \subseteq V \textit{ is an } \overline{s}t\textit{-set}\}.$$

Theorem 1.2 (Menger). *Let $D = (V, A)$ be a digraph, and $s, t \in V$ distinct nodes. The maximum number of edge-disjoint directed paths from s to t is*

$$\min\{\varrho_D(X) : \ X \subseteq V \textit{ is an } \overline{s}t\textit{-set}\}.$$

We use the notation $\lambda_G(s, t)$ for the maximum number of edge-disjoint paths between s and t in G, and $\lambda_D(s, t)$ for the maximum number of edge-disjoint directed paths from s to t in D. These values are called the *local edge-connectivity* between s and t (from s to t).

The fact that the global edge-connectivity of a graph is high can be formulated in different ways:

Proposition 1.3. *For a graph $G = (V, E)$ and a positive integer k, the following are equivalent:*

(i) $\lambda_G(u, v) \geq k$ *for every pair* $u, v \in V$ *of distinct nodes.*

(ii) $d_G(X) \geq k$ *holds for every non-empty proper subset* X *of* V*.*

(iii) To dismantle the graph into 2 components, one needs to delete at least k *edges.*

(iv) The graph remains connected if we delete $k - 1$ *edges.*

A graph G is called *k-edge-connected* (in short, *k-ec*) if the above hold for G. Thanks to the equivalent characterizations, there are different perspectives from which the edge-connectivity of a graph can be viewed, and these will lead us later to different kinds of extensions of the notion.

Sometimes edge-disjoint paths are required only between nodes of a specified subset. Given $S \subseteq V$, a graph G is called *k-edge-connected in S* if there are at least k edge-disjoint paths in G between any two distinct nodes in S. We say that a set X *separates* a set S if $S \cap X \neq \emptyset$ and $S - X \neq \emptyset$. It follows from Theorem 1.1 that a graph G is k-edge-connected in S if and only if $d_G(X) \geq k$ for every $X \subseteq V$ that separates S.

The k-edge-connectivity of digraphs can be defined along similar lines as the undirected case.

Proposition 1.4. *For a digraph $D = (V, A)$ and a positive integer k, the following are equivalent:*

(i) $\lambda_D(u, v) \geq k$ *for every pair* $u, v \in V$ *of distinct nodes.*

(ii) $\varrho_D(X) \geq k$ *holds for every non-empty proper subset* X *of* V*.*

(iii) The digraph remains strongly connected if we delete $k - 1$ *edges.*

A digraph D is called *k-edge-connected* if it has the above properties. Given $S, T \subseteq V$, D is called *k-edge-connected from S to T* if $\lambda_D(s, t) \geq k$ for every distinct $s \in S$ and $t \in T$.

1.2.3 Trees and arborescences

Some special subgraphs of graphs and digraphs have an obvious relation to connectivity. We consider trees contained in graphs, and arborescences contained in digraphs, where a digraph $D' = (V, A')$ is called an *arborescence* rooted at s if it is a directed tree and $\varrho_{D'}(v) = 1$ for every $v \in V - s$. It is obvious that a graph is connected if and only if it contains a spanning tree, and a digraph is strongly connected if and only if it contains a spanning arborescence rooted at s for every $s \in V$.

Given a graph $G = (V, E)$ and a subset $W \subseteq V$, a subtree $G' = (V', E')$ of G is called a *Steiner tree* for W if $W \subseteq V'$. A graph is connected in W if and only if it contains a Steiner tree for W.

In [70] Tutte investigated the problem of decomposing a graph into a given number of connected spanning subgraphs, which is equivalent to finding a given number of edge-disjoint spanning trees of G. He proved the following fundamental result:

Theorem 1.5 (Tutte). *An undirected graph $G = (V, E)$ contains k edge-disjoint spanning trees (or G can be decomposed into k connected spanning subgraphs) if and only if*

$$e_G(\mathcal{F}) \geq k(|\mathcal{F}| - 1) \tag{1.1}$$

holds for every partition $\mathcal{F}$ of V into non-empty subsets, where $e_G(\mathcal{F})$ denotes the number of edges connecting distinct members of $\mathcal{F}$.

Note that given a graph $G = (V, E)$ and a subset $W \subseteq V$, it is NP-complete to decide whether G contains k edge-disjoint Steiner trees for W. This problem will be discussed in greater detail in Chapter 5.

The following result of Nash-Williams [63] is in some sense a complementary pair of Tutte's theorem:

Theorem 1.6 (Nash-Williams). *A graph G can be covered by k forests if and only if $i_G(X) \leq k(|X| - 1)$ for every non-empty subset X of V.*

For digraphs, the following result of Edmonds [14] is an analogue of Tutte's theorem:

Theorem 1.7 (Edmonds). *Let $D = (V, A)$ be a digraph, $s \in V$ a fixed root node. Then D contains k edge-disjoint spanning arborescences rooted at s if and only if*

$$\varrho_D(X) \geq k \quad \text{for every } \emptyset \neq X \subseteq V - s.$$

In fact, Edmonds stated the result in a more general form:

Theorem 1.8 (Edmonds). *Let $D = (V, A)$ be a digraph, and $S_1, \ldots, S_k$ subsets of V; for $X \subseteq V$, let $f(X)$ denote the number of sets S_i not disjoint from X. Then D can be decomposed into directed subgraphs $D_1, \ldots, D_k$ such that D_i is connected from S_i if and only if*

$$\varrho_D(X) \geq k - f(X) \quad \text{for every } \emptyset \neq X \subseteq V.$$

1.2.4 Partition-connectivity and rooted connectivity

As it was indicated in the Overview, we consider different variants of edge-connectivity in the thesis. Here we introduce connectivity properties that depend on the edge-disjoint spanning trees and arborescences contained in a (di)graph.

Proposition 1.9. *Given a graph $G = (V, E)$ and a positive integer k, the following are equivalent by Theorem 1.5:*

(i) G can be decomposed into k edge-disjoint connected spanning subgraphs.

(ii) G contains k edge-disjoint spanning trees.

(iii) To dismantle G into $t + 1$ components for some t, one needs to delete at least kt edges.

(iv) $e_G(\mathcal{F}) \geq k(|\mathcal{F}| - 1)$ holds for every partition $\mathcal{F}$ of V.

A graph G is called *k-partition-connected* if it has the above properties. Note that G is 1-partition-connected if and only if it is connected. A k-partition-connected graph is always k-edge-connected, but the converse is generally not true.

The following is a common generalization of edge-connectivity and partition-connectivity. A graph is called *(k, l)-partition-connected* for positive integers k, l if it remains k-partition-connected after the deletion of any l edges. Equivalently, $e_G(\mathcal{F}) \geq k(|\mathcal{F}| - 1) + l$ holds for every nontrivial partition $\mathcal{F}$ of V. Obviously G is k-partition-connected if and only if it is $(k, 0)$-partition-connected. Simple calculation shows that for $k \leq l$, a graph G is (k, l)-partition-connected if and only if it is $(k + l)$-edge-connected.

As for digraphs, the following is true by Theorem 1.7:

Proposition 1.10. *Given a digraph $D = (V, A)$, a fixed root node s, and a positive integer k, the following are equivalent:*

(i) D can be decomposed into k edge-disjoint spanning sub-digraphs that are connected from s.

(ii) D contains k edge-disjoint spanning arborescences rooted at s.

(iii) There are k edge-disjoint paths from s to any other node.

(iv) $\varrho_D(X) \geq k$ for every non-empty subset X of $V - s$.

A digraph D is called *k-rooted-connected* from root s if it has the above properties. It is easy to see that a graph G has a k-rooted-connected orientation from a given node s if and only if it is k-partition-connected. Indeed, if G is k-partition-connected, then by Proposition 1.9 it contains k edge-disjoint spanning trees, which can be oriented so as to obtain k edge-disjoint arborescences rooted at s. Conversely, if D is a k-rooted-connected orientation of G, then by Proposition 1.10 it contains k edge-disjoint spanning arborescences, thus G contains k edge-disjoint spanning trees.

A common generalization of k-edge-connectivity and k-rooted-connectivity of digraphs can be formulated easily. Given a fixed root node s and positive integers k, l, a digraph D is called *(k, l)-edge-connected* from root s if it contains k edge-disjoint paths from s to any other node, and l edge-disjoint paths to s from any other node. Equivalently, $\varrho_D(X) \geq k$ for every non-empty subset X of $V - s$, and $\varrho_D(X) \geq l$ for every proper subset X of V containing s. Obviously (k, k)-edge-connectivity corresponds to k-edge-connectivity, while $(k, 0)$-edge-connectivity corresponds to k-rooted-connectivity.

It should be noted that a graph G has a (k, l)-edge-connected orientation from a given root s if and only if it has a (k, l)-edge-connected orientation from any root. To see this, suppose that D is a (k, l)-edge-connected orientation of G with root s_1, and s_2 is the desired root. It suffices to see the case $k \geq l$ (since reversing all edges of D switches the role of k and l). By definition there are k edge-disjoint paths from s_1 to s_2. Let us reverse the edges on $k - l$ of these paths. Then $\varrho_D(X)$ decreases by $k - l$ if X is an $\overline{s_1}s_2$-set, it increases by $k - l$ if X is an $\overline{s_2}s_1$-set, and remains unchanged otherwise. So the new orientation is (k, l)-edge-connected from s_2.

Chapter 6 will discuss (k, l)-edge-connected orientations in more detail.

1.3 Hypergraphs and directed hypergraphs

Since the thesis discusses hypergraphs and directed hypergraphs from the perspective of connectivity and orientations, this introduction focuses mainly on these themes.

1.3.1 Connectivity of hypergraphs

Graphs are considered with the possibility of loops and parallel edges, and the same approach is used for hypergraphs. A *hypergraph* is denoted as $H = (V, \mathcal{E})$, where V is the set of nodes and $\mathcal{E}$ is the set of hyperedges. *Hyperedges* are considered to be multisets. To a hyperedge e we associate the characteristic function $\chi_e : V \to \mathbb{Z}_+$, i.e. $\chi_e(v)$ equals the multiplicity of the node v in the hyperedge e. Some special notations will be used when describing the relation of a hyperedge e and a set X:

- $v \in e$ means that $\chi_e(v) > 0$,
- $|e| = \chi_e(V)$,
- $|e \cap X| = \chi_e(X)$,
- $e \subseteq X$ means that $\chi_e(V - X) = 0$,
- The hyperedge $e \cap X$ is defined as $\chi_{e \cap X}(v) := \chi_e(v) * \chi_X(v)$,
- The hyperedge $e - X$ is defined as $e \cap (V - X)$,
- The hyperedge $e_1 - e_2$ is defined as $\chi_{e_1 - e_2}(v) := (\chi_{e_1}(v) - \chi_{e_2}(v))^+$,
- For $v \in V$, the hyperedge $e + v$ is defined as $\chi_{e+v} := \chi_e + \chi_{\{v\}}$,
- For a hyperedge set $\mathcal{E}'$, $\cup(\mathcal{E}')$ is the smallest subset X of V for which $e \subseteq X$ for every $e \in \mathcal{E}'$.

For a positive integer ν, a *ν-hyperedge* is a hyperedge e with $|e| = \nu$. A hypergraph is *ν-uniform* if the cardinality of every hyperedge is ν. By the *rank* of a hypergraph we mean the cardinality of its largest hyperedge. A hyperedge e of $H = (V, \mathcal{E})$ is *induced* by a subset X of V if $e \subseteq X$. The number of hyperedges induced by X is denoted by $i_H(X)$. The *degree* of a node $v \in V$ is $d_H(v) := \sum_{e \in \mathcal{E}} \chi_e(v)$.

In a hypergraph H, a *path* between nodes s and t is an alternating sequence of distinct nodes and hyperedges $s = v_0, e_1, v_1, e_2, \dots, e_k, v_k = t$, such that $v_{i-1}, v_i \in e_i$ for all $i = 1, \dots, k$. Figure 1.1 shows an example of a path between two nodes. H is *connected* if there is a path between any two distinct nodes. A hyperedge e *enters* a set X if $e \cap X \neq \emptyset$ and $e \cap (V - X) \neq \emptyset$. It is easy to see that H is connected if and only if every non-empty proper subset of V is entered by at least one hyperedge of H.

For a hypergraph $H = (V, \mathcal{E})$, we define $\Delta_H(X) = \{e \in \mathcal{E} : e \text{ enters } X\}$, and $d_H(X) := |\Delta_H(X)|$. Note that $d_H(\{v\})$ and $d_H(v)$ can be different, since the hyperedges are multisets. For subsets $X, Y \subseteq V$ let $d_H(X, Y)$ be the number of hyperedges $e \in \mathcal{E}$ with $e \subseteq X \cup Y$, $e \not\subseteq X$, $e \not\subseteq Y$. Every hypergraph has the following properties:

$$d_H(X) + d_H(Y) \geq d_H(X \cap Y) + d_H(X \cup Y) \quad \text{for every } X, Y \subseteq V, \tag{1.2}$$

$$i_H(X) + i_H(Y) = i_H(X \cap Y) + i_H(X \cup Y) - d_H(X, Y) \quad \text{for every } X, Y \subseteq V. \tag{1.3}$$

It is well known that Theorem 1.1 of Menger can be generalized for hypergraphs:

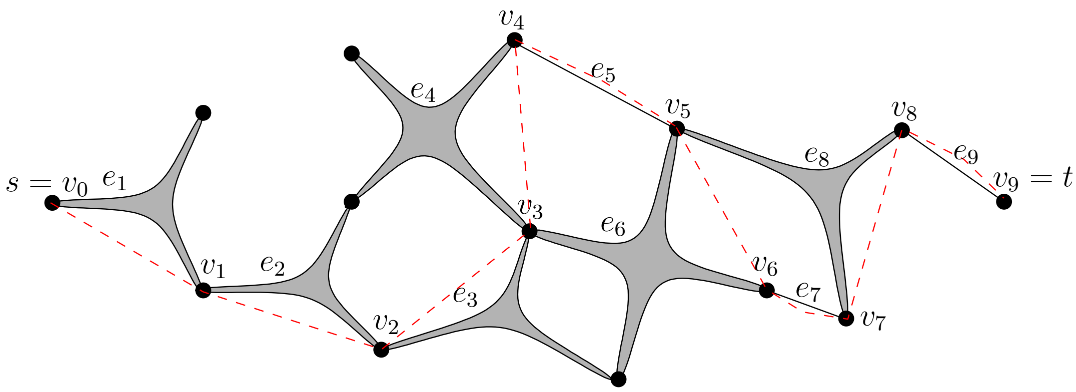

Figure 1.1: A path between s and t in a hypergraph

Theorem 1.11. *Let $H = (V, \mathcal{E})$ be a hypergraph, and $s, t \in V$ distinct nodes. The maximum number of edge-disjoint paths between s and t is*

$$\min\{d_H(X) : \ X \subseteq V \text{ is an } \overline{s}t\text{-set}\}.$$

As for graphs, $\lambda_H(s, t)$ denotes the maximum number of edge-disjoint paths between s and t, and it is called the local edge-connectivity between s and t. For a positive integer k, a hypergraph $H = (V, \mathcal{E})$ is called *k-edge-connected* if the following equivalent conditions hold:

(i) *$\lambda_H(u, v) \geq k$ for every pair $u, v \in V$ of distinct nodes.*

(ii) *$d_H(X) \geq k$ holds for every non-empty proper subset X of V.*

(iii) *To dismantle H into 2 components, one needs to delete at least k hyperedges.*

(iv) *H remains connected if we delete $k - 1$ hyperedges.*

A hypergraph H is said to *cover* a set function p if $d_H(X) \geq p(X)$ for every $X \subseteq V$. So if we define the set function p_k as

$$p_k(X) := \begin{cases} k & \text{if } \emptyset \neq X \subset V, \\ 0 & \text{if } X = \emptyset \text{ or } X = V, \end{cases}$$

then H is k-edge-connected if and only if it covers p_k.

1.3.2 A natural generalization of digraphs: directed hypergraphs

The concept of directed hypergraphs was introduced in many different contexts, in areas like propositional logic, assembly, and relational databases, to efficiently model many-to-

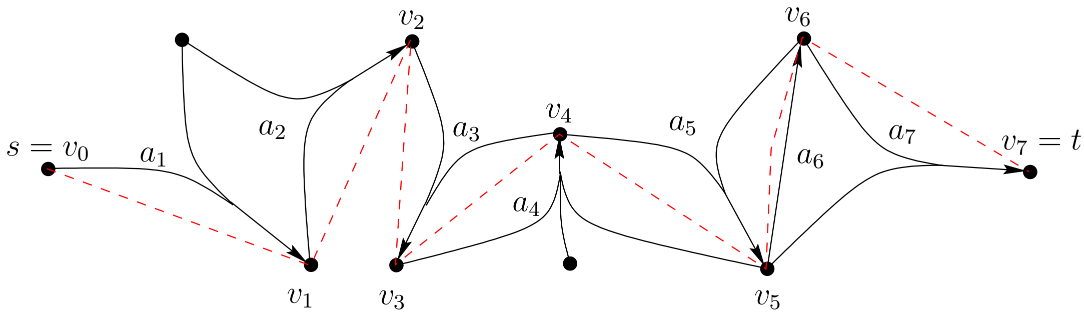

Figure 1.2: A path from s to t in a directed hypergraph

one relations; surveys of these applications can be found in [40] and [41]. There are different possible definitions; we will use the one that best fits our framework for edge-connectivity.

In our terminology, a *directed hypergraph* is a pair $D = (V, \mathcal{A})$, where V is a finite ground set, and $\mathcal{A}$ is a finite collection of so-called hyperarcs (possibly with repetition). A *hyperarc* a is a hyperedge (which we will also denote by a, if this causes no confusion) with a designated *head node* $h(a) \in a$; the rest of its nodes are called *tail nodes*, and are denoted by $t(a)$. So the role of head and tails is asymmetric: while $h(a)$ is a node, $t(a)$ is a multiset. A natural way of looking at a directed hypergraph is that it is an *orientation* of a hypergraph $H = (V, \mathcal{E})$, i.e., a head node $h(e) \in e$ is designated for every hyperedge $e \in \mathcal{E}$. The *underlying hypergraph* of a directed hypergraph D is the one obtained by considering each hyperarc as a hyperedge.

A hyperarc a with r tail-nodes is called an $(r, 1)$*-hyperarc*. $D = (V, \mathcal{A})$ is $(r, 1)$*-uniform* if every hyperarc is an $(r, 1)$-hyperarc. A hyperarc a of D is *induced* by a subset $X \subseteq V$ if $a \subseteq X$. The number of hyperarcs of D induced by X is denoted by $i_D(X)$. The *in-degree* of a node $v \in V$ is $\varrho_H(v) = |\{a \in \mathcal{A} : \ h(a) = v\}|$. The *out-degree* of $v \in V$ is $\delta_H(v) = \sum_{a \in \mathcal{A}} \chi_{t(a)}(v)$.

A *path* from s to t in a directed hypergraph is an alternating sequence, without repetition, of nodes and hyperarcs $s = v_0, a_1, v_1, a_2, \dots, v_{k-1}, a_k, v_k = t$, where v_{i-1} is one of the tail nodes of a_i, and v_i is the head node of a_i $(i = 1, \dots, k)$. Figure 1.2 shows an example of a path in a directed hypergraph. The node t is said to be *reachable* from the node s if there is a path from s to t. A directed hypergraph D is *strongly connected* if every node is reachable from every other node.

A hyperarc a *enters* a set $X \subseteq V$ if $h(a) \in X$ and $t(a) \not\subseteq X$. For a directed hypergraph $D = (V, \mathcal{A})$ we define $\Delta_D^-(X) = \{a \in \mathcal{A} : \ a \text{ enters } X\}$, $\Delta_D^+(X) = \{a \in \mathcal{A} : a \text{ enters } V - X\}$, $\varrho_D(X) := |\Delta_D^-(X)|$, and $\delta_D(X) = |\Delta_D^+(X)|$. For subsets $X, Y \subseteq V$ let $d_D(X, Y)$ be the number of hyperarcs $a \in \mathcal{A}$ with $a \subseteq X \cup Y$, $a \not\subseteq X$, $a \not\subseteq Y$. The following is true for every directed hypergraph D and subsets $X, Y \subseteq V$:

$$\varrho_D(X) + \varrho_D(Y) = \varrho_D(X \cap Y) + \varrho_D(X \cup Y) + d_D(X, Y), \tag{1.4}$$
$$\delta_D(X) + \delta_D(Y) = \delta_D(X \cap Y) + \delta_D(X \cup Y) + d_D(V - X, V - Y). \tag{1.5}$$

Theorem 1.2 extends naturally to directed hypergraphs:

Proposition 1.12. *In a directed hypergraph $D = (V, \mathcal{A})$, there exist k edge-disjoint paths from node s to node t if and only if $\varrho_D(X) \geq k$ for every $\overline{s}t$-set X.*

Proof. Suppose that $\varrho_D(X) \geq k$ for every $\overline{s}t$-set $X \subseteq V$. To reduce the problem to the digraph case, a new node v_a is added to V for every hyperarc $a \in \mathcal{A}$, and the hyperarc a is replaced by edges uv_a for every $u \in t(a)$, and an edge $v_a h(a)$; let $D' = (V', A')$ be the obtained digraph. There is a one-to-one correspondence between the paths from s to t in D and the paths from s to t in D', and edge-disjointness is preserved. By Theorem 1.2, the maximum number of edge-disjoint paths from s to t is

$$\min\{\varrho_{D'}(X') : \ X' \text{ is an } \overline{s}t\text{-set in } V'\}.$$

For such an X', let $X := X' \cap V$; then $k \leq \varrho_D(X) \leq \varrho_{D'}(X')$. □

As a consequence, local edge-connectivity can be defined similarly as for digraphs: for distinct nodes $s, t \in V$, $\lambda_D(s, t)$ is the maximal number of edge-disjoint paths from s to t. The following is true on global connectivity:

Proposition 1.13. *For a directed hypergraph $D = (V, \mathcal{A})$ and a positive integer k, the following are equivalent:*

(i) $\lambda_D(u, v) \geq k$ for every pair $u, v \in V$ of distinct nodes.

(ii) $\varrho_D(X) \geq k$ holds for every non-empty proper subset X of V.

(iii) D remains strongly connected if we delete any $k - 1$ edges.

A directed hypergraph D is called *k-edge-connected* if the above hold for D. Given $S, T \subseteq V$, D is called *k-edge-connected from S to T* if $\lambda_D(s, t) \geq k$ for every distinct $s \in S$ and $t \in T$.

Like Menger's theorem, Theorem 1.8 of Edmonds can be easily extended to directed hypergraphs. Given a set $S \subseteq V$, a directed hypergraph $D = (V, \mathcal{A})$ is *connected from S* if every node $v \in V$ is reachable from some $s \in S$.

Proposition 1.14 ([36]). *Let $D = (V, \mathcal{A})$ be a directed hypergraph, and $S_1, \dots, S_k$ non-empty subsets of V. For $X \subseteq V$, let $f(X)$ denote the number of sets S_i not disjoint from X. Then D can be decomposed into directed sub-hypergraphs $D_1, \dots, D_k$ such that D_i is connected from S_i if and only if*

$$\varrho_D(X) \geq k - f(X) \quad \textit{for every } \emptyset \neq X \subseteq V.$$

Proof. We prove the theorem by induction on the number of hyperarcs of size at least 3. If every hyperarc is a digraph edge, then we can use Theorem 1.8. Suppose that there is a hyperarc $a \in \mathcal{A}$ with $|a| > 2$. Call a set $\emptyset \neq X \subseteq V$ *tight* if $\varrho_D(X) = k - f(X)$. Note that

$$f(X) + f(Y) \geq f(X \cap Y) + f(X \cup Y),$$

for every $X, Y \subseteq V$, so (1.4) implies that if $X \cap Y \neq \emptyset$ then the intersection and union of tight sets is tight. Let $\mathcal{F}$ be the family of tight sets entered by a. If $\mathcal{F} = \emptyset$ or $\mathcal{F}$ has a unique maximal element X, then we can replace the hyperarc a by a directed edge $uh(a)$ where u is an arbitrary node in $a - X$, and use induction. If $\mathcal{F}$ has at least two maximal elements, say X and Y, then a cannot enter $X \cup Y$, since the union would also be tight, which would contradict the maximality. But then $d_D(X, Y) \geq 1$, so by (1.4) $\varrho_D(X \cap Y) + \varrho_D(X \cup Y) = \varrho_D(X) + \varrho_D(Y) - d_D(X, Y) < 2k - f(X \cap Y) - f(X \cup Y)$, so $X \cap Y$ or $X \cup Y$ would violate the condition. □

As a consequence, the following are equivalent for a directed hypergraph $(D = V, \mathcal{A})$ and a fixed root node $s \in V$:

(i) D can be decomposed into k edge-disjoint spanning directed sub-hypergraphs that are connected from s.

(ii) There are k edge-disjoint paths from s to any other node.

(iii) $\varrho_D(X) \geq k$ for every non-empty subset X of $V - s$.

A directed hypergraph D is called *k-rooted-connected* from root s if it has the above properties.

A directed hypergraph D is said to *cover* a set function p if $\varrho_D(X) \geq p(X)$ for every $X \subseteq V$. Note that k-edge-connectivity and k-rooted-connectivity of directed hypergraphs can be described as the covering of appropriate set functions.

1.4 Families of sets

Most of the min-max theorems presented in the thesis contain conditions involving special kinds of families of sets, and the proofs depend heavily on the manipulation of these families. Some basic notations are presented here, while families with special structures will be discussed in Chapter 2.

1.4.1 Basic definitions

Given a finite ground set V, a *family* of sets is a collection of (not necessarily distinct) subsets of V. The empty family is denoted by $\emptyset$. An example is a partition of a set X which is a family of pairwise disjoint sets whose union is X. A *subpartition* of X is a partition of some $X' \subseteq X$; a partition of V is sometimes simply called a partition. For a family $\mathcal{F}$, we use the notation $\mathrm{co}(\mathcal{F}) := \{V - X : \ X \in \mathcal{F}\}$. If $\mathcal{F}$ is a partition, then $\mathrm{co}(\mathcal{F})$ is called a *co-partition*.

The multiplicities of the sets in a family are always taken into account unless otherwise noted; for example, for a set function p and a family $\mathcal{F}$, $\sum_{X\in\mathcal{F}} p(X)$ counts the value of each set as many times as its multiplicity in $\mathcal{F}$. To a family $\mathcal{F}$ we associate the characteristic function $\chi_{\mathcal{F}} : 2^V \to \mathbb{Z}_+$, i.e. $\chi_{\mathcal{F}}(X)$ equals the multiplicity of the set X in $\mathcal{F}$.

The notation $\mathcal{F}_1 \subseteq \mathcal{F}_2$ is used for $\chi_{\mathcal{F}_1} \leq \chi_{\mathcal{F}_2}$. The *sum* of two families $\mathcal{F}_1$ and $\mathcal{F}_2$, denoted by $\mathcal{F}_1 + \mathcal{F}_2$, is the family with characteristic function $\chi_{\mathcal{F}_1} + \chi_{\mathcal{F}_2}$.

Let $H = (V, \mathcal{E})$ be a hypergraph, and $\mathcal{F}$ a family of subsets of V. We define

$$e_H(\mathcal{F}) := \sum_{e\in\mathcal{E}} \max_{u\in e} |\{X \in \mathcal{F} : \ u \in X, \ e \not\subseteq X\}|. \tag{1.6}$$

It is easy to see that

$$e_H(\mathcal{F}) = \max\left\{ \sum_{X\in\mathcal{F}} \varrho_{\vec{H}}(X) : \ \vec{H} \text{ is an orientation of } H \right\}. \tag{1.7}$$

If $\mathcal{F}$ is a partition, then $e_H(\mathcal{F})$ is the number of hyperedges that are not induced by any member of the partition (these are called *cross-hyperedges*).

1.4.2 Duality

One area where families of sets appear naturally is the theory of duality. We do not give a proper introduction to duality here, just mention a few facts and special formulations that will be needed later.

Let $M \in \mathbb{Q}^{m\times n}$ be a matrix, $p \in \mathbb{Q}^m$, $\beta \in \mathbb{Q}$, and $f, g \in \mathbb{Q}^n$ for which $f \le g$. Consider the LP system

$$\{x \in \mathbb{Q}^n : \; Mx \ge p, \; \mathbf{1}x \le \beta, \; f \le x \le g\}, \tag{1.8}$$

where $\mathbf{1}$ is the all-1 vector. Let $c \in \mathbb{Q}^n$ be a cost vector. According to the duality theorem, if the cost cx of vectors x satisfying (1.8) is bounded from below, then

$$\min\{cx : \; Mx \ge p, \; \mathbf{1}x \le \beta, \; f \le x \le g\} = \\ = \max\{yp - \mu\beta + z_+ f - z_- g : \; yM - \mu\mathbf{1} + z_+ - z_- = c, \; y, \mu, z_+, z_- \ge 0\}. \tag{1.9}$$

The system (1.8) is called *totally dual integral* (TDI for short) if for every integer-valued cost vector c for which the above minimum/maximum exists, there is an integer-valued optimal dual solution (y, μ, z_+, z_-). By the fundamental result of Edmonds and Giles [17], a TDI system has an integer-valued optimal primal solution x for every cost vector $c \in \mathbb{Q}^n$ for which the optimum is bounded.

If M is a $0-1$ matrix, then the rows of M can be considered as subsets $Z_1, \dots, Z_m$ of an n-element ground set V, and we can define the set function $p : 2^V \to \mathbb{Q} \cup \{-\infty\}$ as $p(Z_i) := p_i$ $(i = 1, \dots, m)$ and $p(X) := -\infty$ otherwise. So in this case (1.9) can be written in the following form:

$$\min\left\{\sum_{v\in V} c(v)x(v) : \; x(Z) \ge p(Z) \; \forall Z \subseteq V, \; x(V) \le \beta, \; f \le x \le g\right\} = \\ = \max\left\{\sum_{Z\subseteq V} y(Z)p(Z) - \mu\beta + z_+ f - z_- g : \right. \\ \left. \sum_{Z\subseteq V} y(Z)\chi_Z(v) - \mu + z_+(v) - z_-(v) = c(v) \; \forall v \in V, \; y, \mu, z_+, z_- \ge 0\right\}. \tag{1.10}$$

If every value of y is integer (this may be assumed if the system is TDI and c is integer), then y can be considered as the characteristic function of a family of subsets of V.

The feasibility of a linear system is characterized by Farkas' Lemma:

Lemma 1.15. *Given a matrix $M \in \mathbb{Q}^{m\times n}$, $p \in \mathbb{Q}^m$, $\beta \in \mathbb{Q}$, and $f, g \in \mathbb{Q}^n$ for which $f \le g$, the system*

$$\{x :\in \mathbb{Q}^n : \; Mx \ge p, \; \mathbf{1}x \le \beta, \; f \le x \le g\}$$

is solvable if and only if the system

$$\{y \in \mathbb{Q}^m_+, \; \mu \in \mathbb{Q}_+, \; z_+ \in \mathbb{Q}^n_+, \; z_- :\in \mathbb{Q}^n_+ : \\ yM - \mu\mathbf{1} + z_+ - z_- = 0, \; yp + z_+ f - z_- g > \mu\beta\}$$

has no solution.

Again, if M is 0-1 valued and y is integral, then it can be regarded as a characteristic function of a family of sets. Let us cite another version of Farkas' Lemma that will be used later. For a vector z, $z \gg 0$ means that every coordinate of z is strictly greater than 0, while $z > 0$ means that $z \geq 0$ and at least one coordinate is strictly greater than 0.

Lemma 1.16. *Let $M \in \mathbb{Q}^{m\times n}$ be a matrix. The system*

$$\{x \in \mathbb{Q}^n : \; Mx \gg 0, \; x \geq 0\}$$

is solvable if and only if the system

$$\{y \in \mathbb{Q}^m : \; yM \leq 0, \; y > 0\}$$

has no solution.

1.4.3 Compositions

Most of the families that appear in the thesis are so-called compositions. A family $\mathcal{F}$ of sets is a *composition* of a set $X \subseteq V$ if the value $\sum_{Z\in\mathcal{F}} \chi_Z(v) - \chi_X(v)$ is the same for every $v \in V$. If $X \cap Y = \emptyset$, $\mathcal{F}_1$ is a composition of X, and $\mathcal{F}_2$ is a composition of Y, then $\mathcal{F}_1 + \mathcal{F}_2$ is a composition of $X \cup Y$. A composition of V is called a *regular family*; note that these are also compositions of $\emptyset$. If $\mathcal{F}$ is regular, then so is $\mathrm{co}(\mathcal{F})$.

For a set $X \subseteq V$ and a family $\mathcal{F}$ that is a composition of X, we define the *height* of $\mathcal{F}$ with respect to X:

$$h_X(\mathcal{F}) := \sum_{Z\in\mathcal{F}} \chi_Z(v) - \chi_X(v) \quad \text{for an arbitrary } v \in V. \tag{1.11}$$

We will omit the subscript X if it is unambiguous; in fact, the only ambiguity is that if $\mathcal{F}$ is a regular family, then $h_V(\mathcal{F}) = h_\emptyset(\mathcal{F}) - 1$. Note also that if $\mathcal{F}$ is a partition, then $h_V(\mathcal{F}) = 0$, if it is a co-partition then $h_V(\mathcal{F}) = |\mathcal{F}| - 2$ and if $\mathcal{F} = \emptyset$, then $h_V(\mathcal{F}) = -1$!

If $H = (V, \mathcal{E})$ is a hypergraph and $\mathcal{F}$ is a regular family, then the value $e_H(\mathcal{F})$ defined in (1.6) can be expressed in a simpler form:

$$e_H(\mathcal{F}) = h_\emptyset(\mathcal{F})|\mathcal{E}| - \sum_{X\in\mathcal{F}} i_H(X). \tag{1.12}$$

Moreover, the following holds for regular families:

Proposition 1.17. *If $\mathcal{F}_1$ and $\mathcal{F}_2$ are regular families, then $e_H(\mathcal{F}_1+\mathcal{F}_2) = e_H(\mathcal{F}_1)+e_H(\mathcal{F}_2)$.*

It was already mentioned that if $\mathcal{F}$ is a partition, then $e_H(\mathcal{F})$ is the number of cross-hyperedges. It should be noted that for a co-partition $\mathcal{F}$ it can be expressed as:

$$e_H(\mathcal{F}) = \sum_{e \in H} (|\{X \in \mathcal{F} : \ e \cap X \neq \emptyset\}| - 1). \tag{1.13}$$

A generalization of the above definition of composition will also be used. For a function $c : V \to \mathbb{Z}_+$, a *c-composition* is a family $\mathcal{F}$ for which $\sum_{Z \in \mathcal{F}} \chi_Z(v) - c(v)$ is the same for every $v \in V$. For example in (1.10) if $z_+ \equiv z_- \equiv 0$ then an integral y defines a c-composition. Note that a composition of a set X is a χ_X-composition. For $X, Y \subseteq V$, a $(\chi_X + \chi_Y)$-composition will be called an (X, Y)-*composition*.

Chapter 2

Submodular functions

The solution of many edge-connectivity problems relies on the submodular-type properties of certain set functions. In this chapter the basic notions related to submodularity are introduced, and we show how submodular functions can be constructed from set functions with weaker properties. Section 2.3 contains some new results in this respect (Theorem 2.17), as well as a useful extension of polymatroid intersection (Theorem 2.24). At the end of the chapter a brief account is given of submodular flows and algorithms related to submodularity.

2.1 Laminar and cross-free families

2.1.1 Properties

Let V be a finite ground set. Two sets $X, Y \subseteq V$ are called *co-disjoint* if $X \cup Y = V$. They are *intersecting* if none of $X - Y$, $Y - X$, and $X \cap Y$ is empty; they are called *crossing* if they are intersecting and not co-disjoint.

A family $\mathcal{F}$ of subsets of V is called *laminar* if it contains no intersecting members. $\mathcal{F}$ is *cross-free* if it contains no crossing members (so every laminar family is cross-free). Clearly if $\mathcal{F}$ is cross-free, then so is $\text{co}(\mathcal{F})$. A partition is laminar, and a co-partition is cross-free. The following is a well-known property of regular cross-free families:

Proposition 2.1. *Every regular cross-free family decomposes into partitions and co-partitions.* □

Every cross-free family $\mathcal{F}$ has a *tree-representation* (T, φ), where $T = (W, A)$ is a directed tree, and $\varphi : V \to W$ is a mapping such that $\{\varphi^{-1}(W_a) : \ a \in A\} = \mathcal{F}$, where W_a is the component of $T - a$ entered by a. The tree-representation of a laminar family is an arborescence. Figure 2.1 shows a tree-representation of a cross-free family.

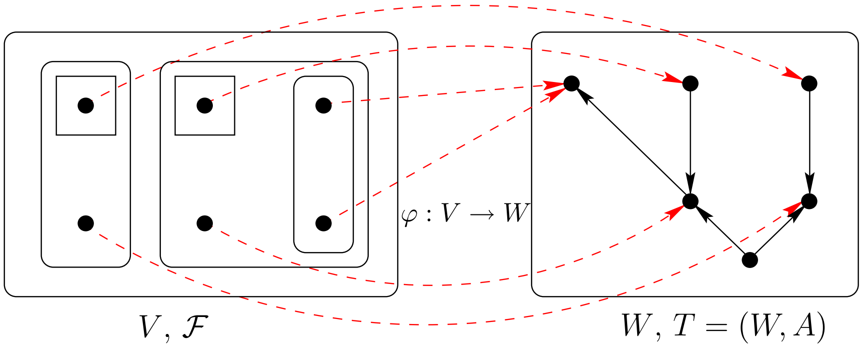

Figure 2.1: Tree-representation of a cross-free family $\mathcal{F}$ (the non-rounded rectangles represent their complement)

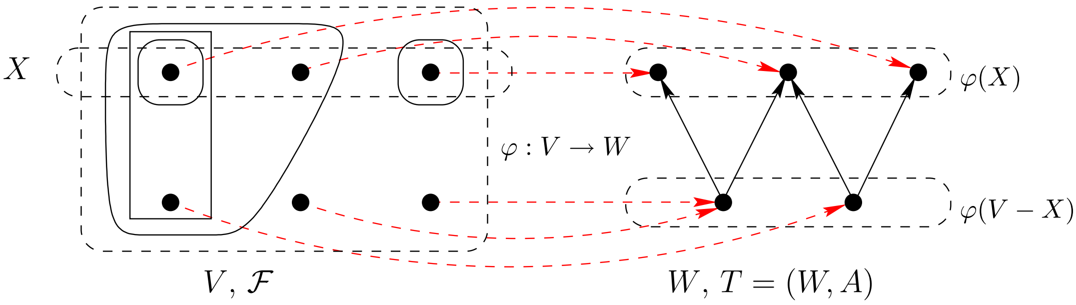

Figure 2.2: A tree-composition $\mathcal{F}$ of X and its tree-representation (the non-rounded rectangles represent their complement)

A *tree-composition* $\mathcal{F}$ of a set $X \subseteq V$ is a cross-free composition of X which contains no partition or co-partition as a proper sub-family (so $\mathcal{F}$ itself can be a partition or a co-partition, and it can also be empty). The name comes from the property that if $\emptyset \neq X \subset V$, then $\mathcal{F}$ has a tree-representation $(T = (W, A), \varphi)$ such that $\varphi^{-1}(w) \neq \emptyset$ for every $w \in W$ (as on Figure 2.2). In this tree-representation, $\varrho_T(w) = 0$ if $\varphi^{-1}(w) \subseteq V - X$, and $\delta_T(w) = 0$ if $\varphi^{-1}(w) \subseteq X$. This implies the following, if we consider a node of W of the latter type that is entered by more than one edge of T:

Claim 2.2. *If $\mathcal{F} \neq \emptyset$ is a tree-composition of X that is not a partition of X, then it contains a subfamily $\{Z_1, \dots, Z_t\}$ $(t \geq 2)$ of pairwise co-disjoint sets such that $\cap Z_i \subseteq X$. If $X \neq V$, then $Z_i - X \neq \emptyset$ $(i = 1, \dots, t)$.* □

Cross-free families can be decomposed into tree-compositions:

Lemma 2.3. *Let the function $c : V \to \mathbb{Z}_+$ have maximum value k, and let $\mathcal{F}$ be a cross-free c-composition. Let $Z_i := \{v \in V : c(v) \geq i\}$ $(i = 1, \dots, k)$. Then $\mathcal{F}$ decomposes* *into partitions, co-partitions, and (possibly empty) families $\mathcal{F}_1, \dots, \mathcal{F}_k$, where $\mathcal{F}_i$ is a tree-composition of Z_i ($i = 1, \dots, k$).*

Proof. If $k = 1$, then $\mathcal{F}$ is a partition of a set $X \subseteq V$, which is a tree-composition. We prove the Lemma by induction on k. It can be assumed that $\mathcal{F}$ contains no partitions and co-partitions of V, since in that case we can use induction. Consider a tree-representation $(T = (W, A), \varphi)$ of $\mathcal{F}$. Recall that for $a \in A$, W_a denotes the component of $T - a$ entered by a. Let $W_i := \{w \in W : |\{a \in A : w \in W_a\}| = i\}$. Then the following are true:

- $Z_i = \cup_{j=i}^{k} \varphi^{-1}(W_i)$,
- If the head of an edge $a \in A$ is in W_i, then its tail is in W_{i-1}.

Let l be the maximal value for which $W_l \neq \emptyset$. If $w \in W_l$, then $\varphi^{-1}(w) \neq \emptyset$, since otherwise $\{\varphi^{-1}(W_a) : a \text{ enters } w\}$ would be a co-partition. Hence $l = k$, and $\varphi^{-1}(W_k) = Z_k$. It is easy to see that for $w \in W_k$, $\{\varphi^{-1}(W_a) : a \text{ enters } w\}$ is a composition of $\varphi^{-1}(w)$. Thus $\mathcal{F}_k := \{\varphi^{-1}(W_a) : a \text{ enters } W_k\} = \cup_{w \in W_k} \{\varphi^{-1}(W_a) : a \text{ enters } w\}$ is a composition of Z_k. It is furthermore cross-free and contains no partitions and co-partitions, so it is a tree-composition of Z_k. We can remove $\mathcal{F}_k$ from $\mathcal{F}$ and use induction on the remaining family. □

If a non-empty tree-composition of $X \subseteq V$ is laminar, then it is a partition of X; thus in case of a laminar family $\mathcal{F}$ Lemma 2.3 gives a decomposition of $\mathcal{F}$ into partitions of $Z_1, \dots, Z_k$.

If $\mathcal{F}$ is a cross-free (X, Y)-composition for some $X, Y \subseteq V$, then Lemma 2.3 implies that it decomposes into a composition of $X \cap Y$ and a composition of $X \cup Y$.

2.1.2 Uncrossing

Uncrossing is a general name for methods which transform a family into a cross-free family by the repeated application of some easy steps. These will be used many times in the thesis; here we only describe the most basic uncrossing method.

Let $y : 2^V \to \mathbb{Q}_+$ be a non-negative set function. By the *uncrossing operation* we mean the following modification of y: given two crossing sets X_1 and X_2 with $y(X_1), y(X_2) > 0$, decrease $y(X_1)$ and $y(X_2)$ by $\min\{y(X_1), y(X_2)\}$, and increase $y(X_1 \cap X_2)$ and $y(X_1 \cup X_2)$ by the same amount. If $y(X)$ is defined as the multiplicity of X in a family $\mathcal{F}$, then we speak of *uncrossing* $\mathcal{F}$.

Lemma 2.4. *After finitely many uncrossing operations y is positive only on a cross-free family of sets.*

Proof. This well-known result can be seen as a special case of the following claim (note that the claim does not hold for non-negative real numbers!):

Claim 2.5. *Let $x_1, \ldots, x_n$ be non-negative rational numbers. Suppose that we apply repeatedly the following operation: for some indices $i < j < k < l$ where x_j and x_k are positive, decrease x_j and x_k by $\min\{x_j, x_k\}$, and increase x_i and x_l by $\min\{x_j, x_k\}$. Then this operation can be repeated only a finite number of times.*

Proof. By multiplying all x_i values by a suitable integer, we can assume that every x_i is integer. Now suppose that there is an infinite sequence of operations, and let m be the smallest index for which x_m decreases infinitely many times. Then one of $x_1, \ldots, x_{m-1}$ increases infinitely many times by at least 1, but decreases finitely many times, which is impossible since $\sum_{i=1}^n x_i$ remains constant and $x_i \geq 0$ for every i. □

Let $X_1, \ldots, X_t$ be an ordering of the subsets of V compatible with the standard partial order; let $x_i := y(X_i)$. Then it follows from the claim that after finitely many uncrossing steps uncrossing is impossible, therefore y is positive on a cross-free family. □

To illustrate the usefulness of the uncrossing technique, let $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be a set function for which $p(X) + p(Y) \leq p(X \cap Y) + p(X \cup Y)$ whenever X and Y are crossing. Lemma 2.4 implies that for any $c : V \to \mathbb{Z}_+$ and any c-composition $\mathcal{F}$ there is a cross-free c-composition $\mathcal{F}'$ such that $\sum_{X \in \mathcal{F}'} p(X) \geq \sum_{X \in \mathcal{F}} p(X)$.

2.2 Submodular and supermodular set functions

2.2.1 Basic properties

A set function $b : 2^V \to \mathbb{Z} \cup \{\infty\}$ is called *fully submodular* (or submodular for short) if

$$b(X) + b(Y) \leq b(X \cap Y) + b(X \cup Y) \tag{2.1}$$

holds for every $X, Y \subseteq V$. It is clear from the definition that the sum of fully submodular set functions is fully submodular. A *modular set function* (obtained from a function $m : V \to \mathbb{Z}$ by $m(X) := \sum_{v \in X} m(v)$) is always fully submodular. There is an alternative way to characterize submodularity:

Proposition 2.6. *A set function $b : 2^V \to \mathbb{Z} \cup \{\infty\}$ is fully submodular if and only if*

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

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

Some examples of fully submodular set functions were already given in Chapter 1. Inequalities (1.2), (1.4), and (1.5) imply that if H is a hypergraph and D is a directed hypergraph, then the set functions $d_H(X)$, $\varrho_D(X)$, and $\delta_D(X)$ are fully submodular. It is also easy to see that if $H = (V, \mathcal{E})$ is a hypergraph, then we can define a fully submodular set function $b : 2^{\mathcal{E}} \to \mathbb{Z}_+$ by $b(\mathcal{E}') := |\cup(\mathcal{E}')|$ for every $\mathcal{E}' \subseteq \mathcal{E}$.

A set function $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ is called *fully supermodular* if

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

holds for every $X, Y \subseteq V$. In other words, p is fully supermodular if and only if $-p$ is fully submodular. An important example is the set function $i_H(X)$ for an arbitrary hypergraph H.

One of the fundamental properties of submodular set functions is that the greedy algorithm can be used to find a maximum weight vector majorized by the set function. For sake of simplicity we describe this only for finite set functions. Let $b : 2^V \to \mathbb{Z}$ be fully submodular, and let $c : V \to \mathbb{Q}$ be a weight function. The aim is to find an $x : V \to \mathbb{Z}$ that satisfies $x(Z) \leq b(Z)$ for every $Z \subseteq V$ and for which $\sum_{v \in V} c(v)x(v)$ is maximal (if the maximum exists).

The greedy algorithm proceeds as follows. We choose an ordering $v_1, \ldots, v_n$ of V such that the values $c(v_1) \geq c(v_2) \geq \cdots \geq c(v_n)$. Let $Z_0 := \emptyset$, and $Z_i := \{v_1, \ldots, v_i\}$ $(i = 1, \ldots, n)$. We repeat for $i = 1, \ldots, n$ the following:

- *Set* $x(v_i) := b(Z_i) - x(Z_{i-1}) = b(Z_i) - b(Z_{i-1})$.

It is clear that the obtained values $x(v_i)$ are integral. It can be shown that $x(Z) \leq b(Z)$ for every $Z \subseteq V$ and $\sum_{v \in V} c(v)x(v)$ is maximal subject to this condition.

2.2.2 Matroids

A prime example of submodular functions is the rank function of a matroid. We include here a few results of matroid theory that will be used later in some proofs. A matroid is given as $M = (S, \mathcal{I})$, where S is the ground set and $\mathcal{I}$ is the collection of *independent sets*, that satisfy the following axioms:

(i) $\emptyset \in \mathcal{I}$,

(ii) If $X \subseteq Y \in \mathcal{I}$, then $X \in \mathcal{I}$,

(iii) For every $X \subseteq S$, the maximal sets among the independent subsets of X have the same cardinality.

The *rank function* of the matroid M, denoted by r_M, is the set function whose value on X is the cardinality of a maximal independent subset of X. This rank function is monotone increasing (i.e. $r_M(X) \leq r_M(Y)$ if $X \subseteq Y$), fully submodular, and $r_M(X) \leq |X|$ for every $X \subseteq S$. It can be shown that every set function satisfying these properties is the rank function of a matroid.

Given a graph $G = (V, E)$, One can define a matroid $M_G = (S, \mathcal{I})$ by $S := E$ and $\mathcal{I} := \{E' \subseteq E : E' \text{ is a forest}\}$. M_G is called the *circuit matroid* of G. A generalization of this construction to hypergraphs, due to Lorea [54], will be presented in Chapter 5.

A graph $G = (V, E)$ is connected if $r_{M_G}(E) = |V| - 1$. G contains k edge-disjoint spanning trees if M_G has k disjoint independent sets of cardinality $|V| - 1$. As the following theorem of Edmonds asserts, this is also a matroid problem.

Theorem 2.7 (Edmonds [13]). *Let $M_i = (S, \mathcal{I}_i)$ be matroids on a common ground set S for $i = 1, \ldots, k$. Then the family $\mathcal{I}_\Sigma := \{I_1 \cup I_2 \cup \cdots \cup I_k : I_i \in \mathcal{I}_i\}$ forms the family of independent sets of a matroid M_Σ whose rank-function r_Σ is given by the following formula:*

$$r_\Sigma(X) = \min \left\{ \sum_{i=1}^{k} r_i(X') + |X - X'| : \; X' \subseteq X \right\}. \tag{2.3}$$

The matroid M_Σ defined in the theorem is called the *sum* of matroids $M_1, \ldots, M_k$. Theorem 1.5 of Tutte can be easily obtained from this result.

Another result of Edmonds that is of central importance is the matroid intersection theorem:

Theorem 2.8 (Edmonds [15]). *Let $M_1 = (S, \mathcal{I}_1)$ and $M_2 = (S, \mathcal{I}_2)$ be two matroids, with rank functions r_1 and r_2. Then the maximum cardinality of a set in $\mathcal{I}_1 \cap \mathcal{I}_2$ equals*

$$\min_{X \subseteq S} (r_1(X) + r_2(S - X)).$$

2.2.3 Polyhedra associated to set functions

The properties of sub- and supermodular functions can be best described using the terminology of polyhedral combinatorics. Here we formulate the results from the perspective of supermodular functions, since these appear more often in the thesis; analogous statements are of course true for submodular functions. In the following paragraphs we use the terms "min" and "max" to denote $-\infty / +\infty$ if the values are not bounded. Given a set function $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$, we define the polyhedra

$$C(p) := \{x : V \to \mathbb{Q} : \; x(Y) \geq p(Y) \; \forall Y \subseteq V\}, \tag{2.4}$$

$$B(p) := \{x : V \to \mathbb{Q} : \; x(V) = p(V); \; x(Y) \geq p(Y) \; \forall Y \subseteq V\}. \tag{2.5}$$

A polyhedron P is a *contra-polymatroid* if $P = C(p)$ for some monotone increasing fully supermodular function p. P is a *base polyhedron* if $P = B(p)$ for some fully supermodular function p. We will use the term "system defining $B(p)$ (or $C(p)$)" for the linear systems given above.

If p is fully supermodular, then both $C(p)$ and $B(p)$ have integral vertices. Moreover, given a weight function $c : V \to \mathbb{Q}$, a minimum weight vertex of $C(p)$ or $B(p)$ can be found by the greedy algorithm. If c is non-negative with maximum value k, then

$$\min\{cx :\ x \in C(p)\} = \min\{cx :\ x \in B(p)\} = \\ = \sum_{i=1}^{k} p(Z_i) \quad \text{where } Z_i = \{v \in V :\ c(v) \geq i\}. \quad (2.6)$$

It follows that for a polyhedron P, if $P = C(p)$ or $P = B(p)$ for some fully supermodular function p, then p is uniquely determined, since $p(X) = \min\{\sum_{v\in V} \chi_X(v)x(v) :\ x \in P\}$.

The intersection of two contra-polymatroids and of two base polyhedra is also a polyhedron with integral vertices. The non-emptiness of the intersection is characterized by the following theorem of Edmonds:

Theorem 2.9 (Edmonds). *Let $p_1 : 2^V \to \mathbb{Z} \cup \{-\infty\}$ and $p_2 : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be fully supermodular functions. Then $B(p_1) \cap B(p_2)$ is non-empty if and only if $p_1(V) = p_2(V)$ and*

$$p_1(X) + p_2(V - X) \leq p_1(V) \quad \textit{for every } X \subseteq V.$$

The minimum weight element of these polyhedra can be determined in polynomial time (although not with a simple greedy algorithm). The systems $\min\{cx :\ x(Z) \geq \max\{p_1(Z), p_2(Z)\}\ \forall Z \subseteq V\}$ and $\min\{cx :\ x(Z) \geq \max\{p_1(Z), p_2(Z)\}\ \forall Z \subseteq V,\ x(V) = p_1(V)\}$ are TDI if p_1 and p_2 are fully supermodular. So the following characterizations are true:

Theorem 2.10. *Let $p_1 : 2^V \to \mathbb{Z} \cup \{-\infty\}$, $p_2 : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be fully supermodular functions. If $c : V \to \mathbb{Q}_+$ is a non-negative weight function, then*

$$\min\{cx :\ x \in C(p_1) \cap C(p_2)\} = \\ = \max_{c_1,c_2\geq 0,\ c_1+c_2=c} (\min\{c_1x :\ x \in C(p_1)\} + \min\{c_2x :\ x \in C(p_2)\}). \quad (2.7)$$

If $p_1(V) = p_2(V)$ and $c : V \to \mathbb{Q}$ is a weight function, then

$$\min\{cx :\ x \in B(p_1) \cap B(p_2)\} = \\ = \max_{c_1+c_2=c} (\min\{c_1x :\ x \in B(p_1)\} + \min\{c_2x :\ x \in B(p_2)\}). \quad (2.8)$$

Deciding whether the intersection of 3 base polyhedra is empty or not is NP-complete.

2.2.4 Intersecting and crossing supermodularity

For set functions that appear in connectivity problems, the supermodular inequality (2.2) often does not hold for every pair of sets. For example, given a hypergraph H_0 and a positive integer k, let $p(X) := k - d_{H_0}(X)$ if $\emptyset \neq X \subset V$, and $p(\emptyset) := p(V) := 0$. For a hypergraph H, $H_0 + H$ is k-edge-connected if and only if H covers p. The set function p is not necessarily supermodular on non-crossing sets.

Another example is for a given directed hypergraph D_0 with a fixed root node s: let $p(X) := k - \varrho_{D_0}(X)$ is $\emptyset \neq X \subseteq V - s$, and $p(X) := 0$ otherwise. For a directed hypergraph D, $D_0 + D$ is k-rooted-connected from s if and only if D covers p. Here p is supermodular on intersecting sets.

This shows that in many cases the requirement of full supermodularity must be relaxed. A set function $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ is *crossing supermodular* if (2.2) holds whenever X and Y are crossing. The set function p is *intersecting supermodular* if (2.2) holds whenever X and Y are intersecting.

For crossing supermodular functions, the following theorem of Fujishige characterizes the non emptiness of $B(p)$.

Theorem 2.11 (Fujishige [38]). *Let $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be a crossing supermodular function. Then $B(p)$ is nonempty if and only if*

$$\sum_{i=1}^{t} p(X_i) \leq p(V),$$
$$\sum_{i=1}^{t} p(V - X_i) \leq (t-1)p(V)$$

both hold for every partition $\{X_1, X_2, \ldots, X_t\}$ of V. Furthermore, if $B(p)$ is non-empty, then it is a base polyhedron. □

If p is intersecting supermodular, then the non-emptiness condition reduces to the following:

Proposition 2.12. *Let $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be an intersecting supermodular function. Then $B(p)$ is nonempty if and only if*

$$\sum_{i=1}^{t} p(X_i) \leq p(V)$$

holds for every partition $\{X_1, X_2, \ldots, X_t\}$ of V. Furthermore, if $B(p)$ is non-empty, then it is a base polyhedron.

A further relaxation is suggested by noticing that in a problem of covering a set function p, a set X is irrelevant if $p(X) \leq 0$. A set function p is called *positively crossing supermodular* if (2.2) holds whenever $p(X) > 0$, $p(Y) > 0$, and X, Y are crossing.

If a positively crossing supermodular function p is also symmetric (i.e. $p(X) = p(V - X)$ for every $X \subseteq V$), then

$$p(X) + p(Y) \leq p(X - Y) + p(Y - X) \tag{2.9}$$

holds for every pair (X, Y) for which $p(X) > 0$, $p(Y) > 0$, and $X - Y$ and $Y - X$ are nonempty. This is the case for example for the set function $p(X) = (k - d_H(X))^+$ $(\emptyset \neq X \subset V)$ for a hypergraph H.

The following version of Theorem 2.11 is true for positively crossing supermodular functions:

Theorem 2.13. *Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular function. Then $B(p)$ is nonempty if and only if*

$$\sum_{i=1}^{t} p(X_i) \leq p(V),$$
$$\sum_{i=1}^{t} p(V - X_i) \leq (t-1)p(V)$$

both hold for every partition $\{X_1, X_2, \dots, X_t\}$ of V. Furthermore, if $B(p)$ is non-empty, then it is a base polyhedron. □

2.2.5 Truncations

Theorem 2.11 stated that if p is intersecting or crossing supermodular, and $B(p)$ is nonempty, then $B(p)$ is a base polyhedron, thus there is a fully supermodular function p^* such that $B(p) = B(p^*)$. But how can we explicitly construct p^*? The answer is the operation called truncation.

The *upper truncation* (or upper Dilworth truncation) of a set function $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ is

$$p^\wedge(X) := \max \left\{ \sum_{Z \in \mathcal{F}} p(Z) : \mathcal{F} \text{ is a partition of } X \right\} . \tag{2.10}$$

It is easy to see that $C(p) = C(p^\wedge)$ and $B(p) = B(p^\wedge)$.

Proposition 2.14 (Lovász [57]). *If p is intersecting supermodular, then $p^\wedge$ is fully supermodular. If p is crossing supermodular, then so is $p^\wedge$.*

The *full truncation* of a set function $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ is

$$p^{\uparrow}(X) := \max\left\{\sum_{Z\in\mathcal{F}} p(Z) - h_X(\mathcal{F})p(V) : \ \mathcal{F} \text{ is a tree-composition of } X\right\} . \tag{2.11}$$

Here $B(p) = B(p^{\uparrow})$, and the following is true:

Proposition 2.15 (Frank [25]). *If p is crossing supermodular, and $B(p)$ is non-empty, then $p^{\uparrow}$ is fully supermodular.*

The algorithms to calculate the upper truncation and the full truncation will be briefly discussed in Section 2.5. Another type of truncation appears when we define a matroid using an intersecting submodular set function. Edmonds [15] proved the following theorem:

Theorem 2.16 (Edmonds [15]). *Let $b : 2^S \to \mathbb{Z}_+$ be a non-negative, integer-valued, intersecting submodular set-function. Then*

$$\mathcal{I}_b := \{X \subseteq S : \ b(Y) \geq |Y \cap X| \text{ for every } Y \subseteq S\} \tag{2.12}$$

forms the family of independent sets of a matroid $M_b = (S, \mathcal{I}_b)$ whose rank-function is given by

$$r_{M_b}(X) = \min\left\{\sum_{Z\in\mathcal{F}} b(Z) + |X - (\cup_{Z\in\mathcal{F}} Z)| : \ \mathcal{F} \text{ is a subpartition of } S\right\}. \tag{2.13}$$

Furthermore, if b is monotone increasing, then

$$\mathcal{I}_b = \{X \subseteq S : \ b(Y) \geq |Y| \text{ for every } Y \subseteq X\} \tag{2.14}$$

and

$$r_{M_b}(X) = \min\left\{\sum_{Z\in\mathcal{F}} b(Z) + |X - (\cup_{Z\in\mathcal{F}} Z)| : \ \mathcal{F} \text{ is a subpartition of } X\right\}. \tag{2.15}$$

2.3 Relaxations of supermodularity

In this section we discuss how supermodularity can be further relaxed while some other important properties, like the TDI property of the associated linear system, are retained. Theorem 2.17 is a comprehensive result on possible relaxations that seems to have not been observed before. The other new observation of the section, Theorem 2.24, describes a way to relax the conditions of Theorem 2.9.

2.3.1 Skew supermodular set functions

Let $H_0 = (V, \mathcal{E})$ be a hypergraph, and $r : V^2 \to \mathbb{Z}_+$ a local edge-connectivity requirement, i.e we want to add a hypergraph H to H_0 such that $\lambda_{H_0+H}(u,v) \geq r(u,v)$ for every $u, v \in V$. It is easy to see that $H_0 + H$ has this property if and only if H covers the set function

$$p(X) := \max_{u \notin X,\ v \in X} r(u,v) - d_{H_0}(X).$$

This set function has the following property ($\star$):

($\star$) *Either (2.2) or (2.9) holds for every $X, Y \subseteq V$.*

A set function with the above property is called *skew supermodular.* If $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ is skew supermodular, then the polyhedron $P = \{x : 2^V \to \mathbb{Q} :\ x \geq 0,\ x(Z) \geq p(Z)\ \forall Z \subseteq V\}$ is a contra-polymatroid, whose unique defining fully supermodular function is

$$p'(X) := \max \left\{ \sum_{Z \in \mathcal{F}} p(Z) :\ \mathcal{F} \text{ is a subpartition of } Z \right\}.$$

On the other hand, $B(p)$ is not always a base polyhedron, and its vertices are not necessarily integral.

2.3.2 Truncated supermodularity

We already saw examples where a set function p was not fully supermodular, but $B(p)$ was nevertheless a base polyhedron. In the following paragraphs our aim is to give a characterization of set functions $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ for which the system in the definition of $B(p)$ is a TDI system, and it describes a base polyhedron. For sake of simplicity, we assume without loss of generality that $p(V) = 0$ (the addition of an integer-valued modular set function does not change the properties discussed here). For $X \subseteq V$, let

$$p^*(X) := \max \left\{ \sum_{Z \in \mathcal{F}} p(Z) :\ \mathcal{F} \text{ is a composition of } X \right\}. \tag{2.16}$$

We assume that $p^*(V) = 0$, which is equivalent to the non-emptiness of $B(p)$ by Lemma 1.15. Clearly $B(p) = B(p^*)$, and if the system defining $B(p)$ is TDI, then

$$p^*(X) = \min \left\{ \sum_{v \in V} \chi_X(v) x(v) :\ x \in B(p) \right\}, \tag{2.17}$$

which is fully supermodular if and only if $B(p)$ is a base polyhedron; moreover, if the system defining $B(p^*)$ is TDI, then so is the system defining $B(p)$. Hence p^* is fully supermodular if and only if the system defining $B(p)$ is TDI and it describes a base polyhedron.

We call a set X *significant* if $p(X) = p^*(X)$, and *tight* if $p^*(X) + p^*(V - X) = 0$. We also introduce for $X, Y \subseteq V$ the notation

$$p^*(X, Y) := \max \left\{ \sum_{Z \in \mathcal{F}} p(Z) : \ \mathcal{F} \text{ is an } (X, Y)\text{-composition} \right\}. \tag{2.18}$$

This function clearly has the following properties:

$$p^*(X, Y) \geq \max\{p^*(X) + p^*(Y), p^*(X \cap Y) + p^*(X \cup Y)\}, \tag{2.19}$$

$$p^*(Y) \geq p^*(X, Y) + p^*(V - X). \tag{2.20}$$

The main result of this section is the following theorem:

Theorem 2.17. *If $p(X) + p(Y) \leq p^*(X \cap Y) + p^*(X \cup Y)$ holds whenever X and Y are significant and crossing, then p^* is fully supermodular.*

Observe that if X and Y are not crossing, then $p^*(X) + p^*(Y) \leq p^*(X \cap Y) + p^*(X \cup Y)$ automatically holds. So the condition of the theorem always holds for non-crossing X and Y. Theorem 2.17 implies Proposition 2.15 (it is easy to prove using the uncrossing technique that $p^\uparrow(X) = p^*(X)$ for every $X \subseteq V$ if p is crossing supermodular and $p^*(V) = 0$). Furthermore, it shows that instead of crossing supermodularity, it is sufficient to require $p(X) + p(Y) \leq p^\uparrow(X \cap Y) + p^\uparrow(X \cup Y)$ for every crossing X, Y.

For technical reasons we will prove a theorem that is a bit more general but less elegant than Theorem 2.17. We say that a family $\mathcal{G}$ *generates* p^* if

$$p^*(X) = \max \left\{ \sum_{Z \in \mathcal{F}} p^*(Z) : \ \mathcal{F} \text{ is a composition of } X, \text{ and its members are in } \mathcal{G} \right\}$$

for every $X \subseteq V$. Obviously the family of significant sets generates p^*.

Theorem 2.18. *Suppose that a family $\mathcal{G}$ generates p^*, and $p^*(X, Y) = p^*(X \cap Y) + p^*(X \cup Y)$ holds for every pair $X, Y \in \mathcal{G}$ for which $p^*(X, Y) = p^*(X) + p^*(Y)$. Then p^* is fully supermodular.*

Proof. The main difficulty of the proof is that the composition $\mathcal{F}$ of X that defines $p^*(X)$ in (2.16) can not always be cross-free (an example for this will be presented at the end of the proof). We will show however that the family $\mathcal{F}$ can be assumed to be cross-free restricted to minimal tight sets. First we prove with the help of a few preliminary claims that $p^*(X, Y) = p^*(X \cap Y) + p^*(X \cup Y)$ is true for every tight X and arbitrary Y.

Claim 2.19. *If X is tight, then $p^*(X, Y) = p^*(X) + p^*(Y)$ for every $Y \subseteq V$.*

Proof. $p^*(X,Y) = p^*(X) + p^*(V-X) + p^*(X,Y) \le p^*(X) + p^*(Y)$ by (2.20). □

Claim 2.20. *If X is tight, Y is arbitrary and $p^*(X) + p^*(Y) \le p^*(X\cap Y) + p^*(X\cup Y)$, then $p^*(Y) = p^*(X\cap Y) + p^*(Y-X)$.*

Proof. $p^*(Y) = p^*(Y) + p^*(X) + p^*(V-X) \le p^*(X\cap Y) + p^*(X\cup Y) + p^*(V-X) \le p^*(X\cap Y) + p^*(Y-X)$. □

It follows from the conditions of the theorem and Claims 2.19 and 2.20 that if $X, Y \in \mathcal{G}$ and X is tight, then $p^*(Y) = p^*(X\cap Y) + p^*(Y-X)$.

Claim 2.21. *If $X \in \mathcal{G}$ is tight and Y is arbitrary, then $p^*(X,Y) = p^*(X\cap Y) + p^*(X\cup Y)$.*

Proof. Let $\mathcal{F}_Y$ be a composition of Y whose members are in $\mathcal{G}$ and for which $p^*(Y) = \sum_{Z\in\mathcal{F}_Y} p^*(Z)$ (such a composition exists because $\mathcal{G}$ generates p^*). Claim 2.19 implies that

$$p^*(X,Y) = p^*(X) + p^*(Y) = p^*(X) + \sum_{Z\in\mathcal{F}_Y} p^*(X\cap Z) + \sum_{Z\in\mathcal{F}_Y} p^*(Z-X).$$

There is a non-negative integer ν such that if we add $V-X$ with multiplicity ν to the family $\{X\cap Z : Z\in\mathcal{F}_Y\}$, then we get a composition of $X\cap Y$. It follows that if we add X with multiplicity $\nu+1$ to the family $\{Z-X : Z\in\mathcal{F}_Y\}$, then we get a composition of $X\cup Y$. So

$$p^*(X,Y) = \sum_{Z\in\mathcal{F}_Y} p^*(X\cap Z) + \nu p^*(V-X) + (\nu+1)p^*(X) + \sum_{Z\in\mathcal{F}_Y} p^*(Z-X) \le$$
$$\le p^*(X\cap Y) + p^*(X\cup Y).$$

□

Claim 2.22. *If X is tight and Y arbitrary, then $p^*(X,Y) = p^*(X\cap Y) + p^*(X\cup Y)$.*

Proof. Let $\mathcal{F} = \{X_1,\dots,X_t\}$ be a composition of X with all members in $\mathcal{G}$ such that $p^*(X) = \sum_{i=1}^t p^*(X_i)$. In this case every $X_i \in \mathcal{F}$ is tight, since $p^*(X_i) + p^*(V-X_i) \ge p^*(X_i) + \sum_{j\ne i} p^*(X_j) + p^*(V-X) = 0$. We will construct a cross-free (X,Y)-composition $\mathcal{F}'$ in t steps, for which $p^*(X,Y) = \sum_{Z\in\mathcal{F}'} p^*(Z)$. In every step we have a laminar family $\mathcal{F}_i$, at the beginning $\mathcal{F}_0 := \{Y\}$. In the i-th step with $\mathcal{F}_{i-1} = \{Z_1,\dots,Z_l\}$, let $\mathcal{F}_i := \{Z_1\cap X_i, Z_1 - X_i, Z_2\cap X_i, Z_2 - X_i, \dots, Z_l\cap X_i, Z_l - X_i, X_i\}$. It is easy to see that $\mathcal{F}_i$ is a $(\chi_Y + \sum_{j=1}^i \chi_{X_j})$-composition. It follows from Claims 2.21 and 2.20 that $\sum_{Z\in\mathcal{F}_i} p^*(Z) = p^*(Y) + \sum_{j=1}^i p^*(X_j)$. Let $\mathcal{F}' := \mathcal{F}_t$; then $\mathcal{F}'$ is a laminar (X,Y)-composition, and $\sum_{Z\in\mathcal{F}'} p^*(Z) = p^*(X) + p^*(Y)$. By Lemma 2.3, $\mathcal{F}'$ decomposes into a composition of $X\cap Y$ and a composition of $X\cup Y$; thus $p^*(X,Y) = \sum_{Z\in\mathcal{F}'} p^*(Z) \le p^*(X\cap Y) + p^*(X\cup Y)$. □

According to Claim 2.22 we can assume that all tight sets are in $\mathcal{G}$. If X and Y are tight, then by Claim 2.20, $0 = p^*(Y) + p^*(V - Y) = p^*(X \cap Y) + p^*(Y - X) + p^*(V - Y) \le p^*(X \cap Y) + p^*(V - (X \cap Y)$, so $X \cap Y$ is tight, and similarly $Y - X$ is tight. This means that the minimal tight sets form a partition $W_1, \dots, W_k$. The conditions of the theorem are not changed if we add a modular set function m to p for which $m(V) = 0$. Thus we can assume that $p^*(W_1) = \dots = p^*(W_k) = 0$, so $p^*(X) = 0$ for every tight set X.

We say that sets X, Y are crossing restricted to W_l if $X \cap W_l$ and $Y \cap W_l$ are crossing on the ground set W_l. For a given $1 \le l \le k$ we define a matrix A_l. The rows of A_l correspond to the pairs $X, Y \in \mathcal{G}$ that are crossing restricted to W_l, and for which $p^*(X, Y) = p^*(X) + p^*(Y)$. The columns of A_l correspond to the sets in $\mathcal{G}$. In the following we define the row of A_l corresponding to a pair (X, Y) of sets. There is a composition $\mathcal{F}_{X\cap Y}$ of $X \cap Y$ and a composition $\mathcal{F}_{X\cup Y}$ of $X \cup Y$, both consisting of sets in $\mathcal{G}$, for which $p^*(X, Y) = p^*(X \cap Y) + p^*(X \cup Y) = \sum_{Z\in\mathcal{F}_{X\cap Y}} p^*(Z) + \sum_{Z\in\mathcal{F}_{X\cup Y}} p^*(Z)$. Observe that $\mathcal{F}_{X\cap Y} + \mathcal{F}_{X\cup Y}$ contains at least 2 non-tight sets (because W_l is a minimal tight set), and if it contains exactly 2, then restricted to W_l these are $(X \cap Y) \cap W_l$ and $(X \cup Y) \cap W_l$. Using $\mathcal{F}_{X\cap Y}$ and $\mathcal{F}_{X\cup Y}$, we define the row r of A_l corresponding to X, Y: for a set $Z \in \mathcal{G}$, let $r(Z) := -\chi_{\mathcal{F}_{X\cap Y}}(Z) - \chi_{\mathcal{F}_{X\cup Y}}(Z)$ if $Z \ne X$ and $Z \ne Y$, and let $r(X) = r(Y) := 1 - \chi_{\mathcal{F}_{X\cap Y}} - \chi_{\mathcal{F}_{X\cup Y}}(Z)$.

Claim 2.23. *The system* $\{yA_l \le 0,\ y > 0\}$ *has no solution.*

Proof. Suppose indirectly that there exists such a y; we can assume that y is integral. The vector $-yA_l$ can be considered as a non-negative function whose domain of definition is $\mathcal{G}$. Let $\mathcal{F}$ be the family defined by $\chi_{\mathcal{F}}(Z) := -yA_l(Z)$ $(Z \in A)$, and let $\mathcal{F}'$ be the family obtained by leaving out the tight sets from $\mathcal{F}$. By the definition of the rows of A_l, $\sum_{Z\in\mathcal{F}} p^*(Z) = 0$, and $\mathcal{F}$ is regular. This is possible only if $\mathcal{F}' = \emptyset$. But in a row of A_l corresponding to a pair (X, Y), exactly 2 columns corresponding to non-tight sets have value +1 (those corresponding to X and Y), and at least 2 columns corresponding to non-tight sets have negative value. This is possible only if each row of A_l for which y is positive contains exactly 2 columns corresponding to non-tight sets that have negative value (which is -1). We have seen that for a row corresponding to (X, Y) the intersection of these non-tight sets with W_l must be $(X \cap Y) \cap W_l$ and $(X \cup Y) \cap W_l$. But $|X \cap W_l|^2 + |Y \cap W_l|^2 < |(X \cap Y) \cap W_l|^2 + |(X \cup Y) \cap W_l|^2$, so $\sum_{Z\in\mathcal{F}'} |Z \cap W_l|^2 > 0$, contradicting $\mathcal{F}' = \emptyset$. □

According to the Lemma 1.16 the dual system $\{A_l z \gg 0, z \ge 0\}$ has a solution. Let z_l be an arbitrary solution of the above system (for $l = 1, \dots, k$). This means that z_l is a non-negative weight function on the sets of $\mathcal{G}$, with the following property:

($\star$) *If $X, Y \in \mathcal{G}$ are crossing restricted to W_l, and $p^*(X,Y) = p^*(X) + p^*(Y)$, then there is a composition $\mathcal{F}_{X\cap Y}$ of $X \cap Y$ and a composition $\mathcal{F}_{X\cup Y}$ of $X \cup Y$, both consisting of sets in $\mathcal{G}$, for which $p^*(X,Y) = \sum_{Z\in\mathcal{F}_{X\cap Y}} p^*(Z) + \sum_{Z\in\mathcal{F}_{X\cup Y}} p^*(Z)$ and $z_l(X) + z_l(Y) > \sum_{Z\in\mathcal{F}_{X\cap Y}} z_l(Z) + \sum_{Z\in\mathcal{F}_{X\cup Y}} z_l(Z)$.*

Now we can prove that p^* is fully supermodular. Let $X, Y \subseteq V$. For a given $1 \le l \le k$ let $\mathcal{F}_l$ be an (X,Y)-composition consisting of sets in $\mathcal{G}$, for which $p^*(X,Y) = \sum_{Z\in\mathcal{F}_l} p^*(Z)$, and which is of minimal weight with respect to z_l. (such a family exists, since $\mathcal{G}$ generates p^*). Because of the minimality of the weight, no $X', Y' \in \mathcal{F}_l$ can be crossing restricted to W_l (otherwise X' and Y' could be replaced by the family $\mathcal{F}_{X'\cap Y'} + \mathcal{F}_{X'\cup Y'}$ whose existence is stated in ($\star$), thus decreasing the weight). This means that $\mathcal{F}_l$ is cross-free restricted to W_l.

For $1 \le i \le k$, let $\mathcal{F}_{l,i} := \{Z \cap W_i : Z \in \mathcal{F}_l\}$. On the ground set W_i, $\mathcal{F}_{l,i}$ is an $(X\cap W_i, Y \cap W_i)$-composition. By Claims 2.22 and 2.20, $\sum_{Z\in\mathcal{F}_l} p^*(Z) = \sum_{i=1}^k \sum_{Z\in\mathcal{F}_{l,i}} p^*(Z)$. Since $\sum_{Z\in\mathcal{F}_l} p^*(Z)$ is maximal among (X,Y)-compositions, this implies that $\sum_{Z\in\mathcal{F}_{l_1,i}} p^*(Z) = \sum_{Z\in\mathcal{F}_{l_2,i}} p^*(Z)$ for $1 \le l_1, l_2 \le k$.

On the ground set W_l, the cross-free family $\mathcal{F}_{l,l}$ decomposes into a composition $\mathcal{F}_{l,l}^{\cap}$ of $(X\cap Y) \cap W_l$ and a composition $\mathcal{F}_{l,l}^{\cup}$ of $(X \cup Y) \cap W_l$ by Lemma 2.3. On the ground set V, the family $\sum_{l=1}^k \mathcal{F}_{l,l}^{\cap}$ can be made an $X \cap Y$-composition by the addition of tight sets; let $\mathcal{F}^{\cap}$ denote this $X\cap Y$-composition. Similarly, we can obtain an $X \cup Y$-composition $\mathcal{F}^{\cup}$ from $\sum_{l=1}^k \mathcal{F}_{l,l}^{\cup}$ by the addition of tight sets. Since $p^*(Z) = 0$ for every tight set Z, we get that $p^*(X,Y) = \sum_{Z\in\mathcal{F}^{\cap}} p^*(Z) + \sum_{Z\in\mathcal{F}^{\cup}} p^*(Z)$, hence $p^*(X,Y) \le p^*(X\cap Y) + p^*(X \cup Y)$. This proves the supermodularity of p^*. $\square$

Remark. The following example shows that it is not enough to consider cross-free compositions in the definition of p^*. Let $V = \{v_1, v_2, v_3, v_4\}$ and let $X_1 = \{v_2, v_3, v_4\}$, $X_2 = \{v_1, v_2\}$, and $X_3 = \{v_1, v_3\}$. Let $p(\{v_1\}) = 1$, $p(X_1) = p(X_2) = p(X_3) = -1$, and $p(Z) = -\infty$ on every other set. For the supermodularity of p^* we only need to check that $p(X_2) + p(X_3) \le p^*(X_2 \cap X_3) + p^*(X_2 \cup X_3)$ which is true since $p(X_2 \cap X_3) = 1$ and $p^*(X_2 \cup X_3) \ge p(X_1) + p(X_2) + p(X_3) = -3$. But there is no cross-free $(X_2 \cup X_3)$-composition $\mathcal{F}$ for which $\sum_{Z\in\mathcal{F}} p(Z) \ge -3$.

2.3.3 Jointly supermodular functions

The following relaxation concerns the intersection of base polyhedra, discussed in Theorems 2.9 and 2.10. Let $p_1 : 2^V \to \mathbb{Z} \cup \{-\infty\}$ and $p_2 : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be set functions. By Theorem 2.10, the system defining $B(p_1) \cap B(p_2)$ is TDI if p_1 and p_2 are fully supermodular (and by Proposition 2.15 it is TDI even if p_1 and p_2 are only crossing supermodular).

We consider the relaxation where p_1 is fully supermodular, but for p_2 the supermodular inequality (2.2) is guaranteed to hold only on crossing pairs (X, Y) for which $p_1(X) < p_2(X)$ and $p_1(Y) < p_2(Y)$. This type of set function pairs appear for example in orientation problems featuring positively crossing supermodular set functions, discussed in Chapter 7. The next theorem states that the system defining $B(p_1) \cap B(p_2)$ is TDI even in this relaxed case.

Theorem 2.24. *Let $p_1 : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be fully supermodular, and let $p_2 : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be a set function that is supermodular on the crossing pairs (X, Y) for which $p_1(X) < p_2(X)$ and $p_1(Y) < p_2(Y)$. Then the system*

$$\{x : V \to \mathbb{Q} : \ x(V) = p_1(V); \ x(Z) \geq p_1(Z) \ \forall Z \subseteq V, \ x(Z) \geq p_2(Z) \ \forall Z \subseteq V\} \tag{2.21}$$

is TDI, and it has a feasible solution if and only if

$$p_1(V - X) + \sum_{Z \in \mathcal{F}} p_2(Z) \leq (h_X(\mathcal{F}) + 1)p_1(V) \tag{2.22}$$

for every $X \subseteq V$ and every tree-composition $\mathcal{F}$ of X.

Proof. Let $c : V \to \mathbb{Z}$ be an integral weight function. In the dual system, the variables are $y_1 : 2^V \to \mathbb{Q}_+$, $y_2 : 2^V \to \mathbb{Q}_+$, and $\mu \in \mathbb{Q}_+$. By (1.9), the dual system is

$$\max\left\{\sum_{Z \subseteq V} (y_1(Z)p_1(Z) + y_2(Z)p_2(Z)) - \mu p_1(V) : \right.$$
$$\left. \sum_{Z \subseteq V} (y_1(Z) + y_2(Z))\chi_Z(v) - \mu = c(a) \ \forall v \in V\right\}. \tag{2.23}$$

Suppose that (y_1, y_2, μ) is an optimal dual solution where $\sum_{Z \subseteq V} y_2(Z)$ is minimal. We start the uncrossing procedure on y_2 described in Subsection 2.1.2 (i.e. given two crossing sets X_1 and X_2 with $y_2(X_1), y_2(X_2) > 0$, we decrease $y_2(X_1)$ and $y_2(X_2)$ by $\min\{y_2(X_1), y_2(X_2)\}$, and increase $y_2(X_1 \cap X_2)$ and $y_2(X_1 \cup X_2)$ by the same amount). We claim that all along the uncrossing procedure, y_2 will be positive only on sets Z for which $p_2(Z) > p_1(Z)$. Indeed, if $y_2(Z) > 0$ for a set for which $p_2(Z) \leq p_1(Z)$, then by increasing $y_1(Z)$ by $y_2(Z)$ and decreasing $y_2(Z)$ to 0, we would get an optimal dual solution for which $\sum_{Z \subseteq V} y_2(Z)$ is less.

Let y'_2 be the result of this uncrossing. By Lemma 2.4, y'_2 is positive on a cross-free family $\mathcal{F}_2$ of sets. We can also uncross y_1, and here we can apply the uncrossing operation on non-crossing pairs of sets too, so in the end we obtain y'_1 which is positive on a chain $\mathcal{F}_1$ of sets.

Now (y_1', y_2', μ) is an optimal dual solution of the problem

$$\min\{cx:\ x(V) = p_1(V);\ x(Z) \geq p_1(Z)\ \forall Z \in \mathcal{F}_1,\ x(Z) \geq p_2(Z)\ \forall Z \in \mathcal{F}_2\}. \tag{2.24}$$

This system is the intersection of two base polyhedra, so it has an integral dual optimal solution, which is in turn optimal for (2.23). This proves the TDI property.

The non-emptiness condition can be proved similarly. Suppose that the system 2.21 has no solution. Then by Lemma 1.15, there exists (y_1, y_2, μ) such that

$$\sum_{Z\subseteq V} (y_1(Z) + y_2(Z))\chi_Z(v) = \mu\ \forall v \in V,$$

and

$$\sum_{Z\subseteq V} (y_1(Z)p_1(Z) + y_2(Z)p_2(Z)) - \mu p_1(V) > 0.$$

By the same argument as in the proof of the TDI property, we can assume that y_1 is positive only on the members of a chain $\mathcal{F}_1$ and y_2 is positive only on the members of a cross-free family $\mathcal{F}_2$. But then the system

$$\{x : V \to \mathbb{Q}:\ x(V) = p_1(V);\ x(Z) \geq p_1(Z)\ \forall Z \in \mathcal{F}_1,\ x(Z) \geq p_2(Z)\ \forall Z \in \mathcal{F}_2\}$$

has no solution, so by Theorem 2.9, $p_1(V - X)) + p_2^{\uparrow}(X) > p_1(V)$ for some $X \subseteq V$. □

2.4 Submodular flows

The submodular flow polyhedron of Edmonds and Giles [17] proved to be a very efficient tool in many different areas of combinatorial optimization. Here we just mention a few basic properties, following a brief description of network matrices.

2.4.1 Network matrices

Let $T = (W, A_1)$ be a directed tree, and $D = (W, A_2)$ be a directed graph with the same node set W. We define the network matrix $N = N(W, A_1, A_2)$ using these digraphs. The rows of N are indexed with the edges in A_1, while the columns are indexed with edges in A_2. Let $uv \in A_2$. There is a unique (not necessarily directed) path in T from v to u. We define the column of N corresponding to uv using this path: if $a \in A_1$ is a forward-edge of the path, then the value in the corresponding row is 1; if it is a backward-edge of the path, then the value is -1; if it does not belong to the path, then the value is 0.

It is well known that network matrices are totally unimodular, so the following is true:

Proposition 2.25. *Let $N = N(W, A_1, A_2)$ be a network matrix. Let $l : A_1 \to \mathbb{Z} \cup \{-\infty\}$, $u : A_1 \to \mathbb{Z} \cup \{\infty\}$ be lower and upper capacities on A_1 such that $l \leq u$, and let $f : A_2 \to \mathbb{Z} \cup \{-\infty\}$, $g : A_2 \to \mathbb{Z} \cup \{\infty\}$ be lower and upper capacities on A_2 such that $f \leq g$. Then the system*

$$\{x : A_2 \to \mathbb{Q} : \ l \leq Nx \leq u, \ f \leq x \leq g\} \tag{2.25}$$

is TDI.

To see an example, consider a digraph $D = (V, A)$, and a cross-free family $\mathcal{F}$ on the ground set V. Let $(T = (W, A_1), \varphi)$ be the tree-representation of $\mathcal{F}$. We can associate a digraph $D' = (W, A_2)$ to D by taking the edge $\varphi(u)\varphi(v)$ for every edge $uv \in A$. We get a network matrix $N = N(W, A_1, A_2)$. The system (2.25) described in Proposition 2.25 corresponds to the following system, if we consider l, u to be given on the sets of $\mathcal{F}$, and f, g to be given on the edges of D:

$$\{x : A \to \mathbb{Q} : \ l(Z) \leq \delta_x(Z) - \varrho_x(Z) \leq u(Z) \ \forall Z \in \mathcal{F}, \ f \leq x \leq g\}, \tag{2.26}$$

where $\delta_x(Z) = \sum_{a \in \Delta_D^+(Z)} x(a)$ and $\varrho_x(Z) = \sum_{a \in \Delta_D^-(Z)} x(a)$. So by proposition 2.25, this system is TDI.

2.4.2 Submodular flows

To simplify the notations later in the thesis, we describe submodular flows using supermodular functions, but this is of course equivalent to the submodular formulation.

Let $D = (V, A)$ be a digraph, and $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ a crossing supermodular function. Let furthermore $f : A \to \mathbb{Z} \cup \{-\infty\}$, $g : A \to \mathbb{Z} \cup \{\infty\}$ be upper and lower capacities on A such that $f \leq g$. For $x : A \to \mathbb{Q}$, let $\phi_x(Z) := \delta_x(Z) - \varrho_x(Z)$, which is a modular function. The system

$$\{x : A \to \mathbb{Q} : \ \phi_x(Z) \geq p(Z) \ \forall Z \subseteq V, \ f \leq x \leq g\} \tag{2.27}$$

is called a *submodular flow system.* A simple example is system (2.26), where $p(X) = \max\{l(X), -u(V - X)\}$ if $X \in \mathcal{F} + \text{co}(\mathcal{F})$,and $p(X) = -\infty$ otherwise.

A submodular flow system is *one-way* if $p(X) > -\infty$ implies that either $\varrho_D(X) = 0$ or $\delta_D(X) = 0$. It is *strongly one-way* if either $\varrho_D(X) = 0$ whenever $p(X) > -\infty$, or $\delta_D(X) = 0$ whenever $p(X) > -\infty$.

The fundamental result on submodular flows is the following:

Theorem 2.26 (Edmonds and Giles[17]). *The system (2.27) is TDI.*

Algorithms for finding a minimum cost submodular flow will be mentioned in Section 2.5. The solvability of system (2.27) has a simple characterization if p is fully supermodular:

Theorem 2.27 (Frank [23]). *Suppose that p is fully supermodular. Then (2.27) has a solution if and only if $\delta_g(Z) - \varrho_f(Z) \geq p(Z)$ for every $Z \subseteq V$. If there is a solution, then there is an integral solution as well.*

Suppose now that p is crossing supermodular, and $p(V) = 0$ (this can be assumed since (2.27) has no solution if $p(V) > 0$). Since ϕ_x is modular, the system (2.27) can be written in the form

$$\{x : A \to \mathbb{Q} : \ \phi_x \in B(p), \ f \leq x \leq g\}. \tag{2.28}$$

If $B(p)$ is non-empty, then by Proposition 2.15, the set function $p^\uparrow$ is fully supermodular, and $B(p) = B(p^\uparrow)$. Thus Theorem 2.27 implies the following:

Theorem 2.28 (Frank [23]). *Suppose that p is crossing supermodular. Then (2.27) has a solution if and only if $\delta_g(Z) - \varrho_f(Z) \geq p^\uparrow(Z)$ for every $Z \subseteq V$, where $p^\uparrow$ is the full truncation of p, defined in (2.11). If there is a solution, then there is an integral solution as well.*

2.5 Algorithms

2.5.1 Oracles

When set functions are considered from an algorithmic point of view, the way of getting information about their values must be clarified. Many set functions in the thesis are defined using a graph or a hypergraph, and in these cases the values usually can be computed using network flows. However, in the general case, we consider the values to be given by some kind of oracle. We distinguish two kinds of oracles for a set function $b : 2^V \to \mathbb{Z} \cup \{\infty\}$:

Evaluation oracle: Provides the value of $b(X)$ for any subset $X \subseteq V$. It also tells for $s, t \in V$ whether there is an $\overline{s}t$-set Z with $b(Z) < \infty$.

Minimizing oracle: Given $\emptyset \subseteq X \subseteq Y \subseteq V$ and a modular function m, it calculates $\min\{b(Z) - m(Z) : \ X \subseteq Z \subseteq Y\}$.

A *maximizing oracle* for a set function $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ is a minimizing oracle for $-p$. It was shown in [32] that the upper truncation 2.10 of an intersecting supermodular function p, or the full truncation 2.11 of a crossing supermodular function p can be evaluated in polynomial time (by a combinatorial algorithm), given a maximizing oracle for p. More generally, if there is a maximizing oracle available for a set function p, and the set function p^* is defined by (2.17), then p^* can be evaluated in polynomial time using the ellipsoid method. This implies for example that if p has a maximizing oracle, and the upper truncation (full truncation) of p is fully supermodular, then the upper truncation (full truncation) can be evaluated in polynomial time.

2.5.2 Submodular function minimization

In [43], Grötschel, Lovász, and Schrijver constructed a polynomial time minimizing oracle for any submodular function given by an evaluation oracle, relying on the ellipsoid method. Recently, combinatorial algorithms for submodular function minimization were found independently by Schrijver [67] and by Iwata, Fleicher, and Fujishige [45]. Since then, these algorithms have been further improved.

An algorithm for minimizing fully submodular functions automatically gives an algorithm for minimizing crossing supermodular functions, since for every $s, t \in V$ the set function

$$b_{st}(X) := \begin{cases} b(X) & \text{if } X \text{ is an } \overline{s}t\text{-set}, \\ \infty & \text{otherwise} \end{cases}$$

is fully supermodular. So we can minimize b_{st} for every pair s, t.

The minimum cost submodular flow problem can also be solved in polynomial time. There is an obvious relation with submodular function minimization: checking whether a submodular flow is feasible is a submodular function minimization problem. But the minimum cost submodular flow problem is considered to be much more difficult than submodular function minimization, and current algorithms are much slower (usually involving multiple calls to a submodular function minimization oracle).

It is important to note that for positively crossing supermodular functions in general, no maximization oracle is obtainable based on an evaluation oracle; indeed, an evaluation oracle cannot decide in polynomial time whether there is a set $Z \subseteq V$ for which $p(Z) > 0$.

The maximization of a submodular function is NP-hard.

Chapter 3

Edge-connectivity augmentation of hypergraphs

3.1 Introduction

Of the different types of problems related to edge-connectivity, the one that is described in greatest detail in the thesis is edge-connectivity augmentation. In general, a connectivity augmentation problem consists of optimally augmenting a (hyper)graph or a directed (hyper)graph by adding new edges, so as to meet a specified connectivity requirement. The optimality of the augmentation can be defined in several ways: the added (hyper)edges should satisfy a degree specification on the node set, the number of added edges should be minimized, or the cost of the added edges should be minimized according to some cost function. Usually we will consider the first two types of criteria; in many cases these problems are solvable, while the corresponding minimum cost augmentation is NP-hard.

In this chapter, we describe edge-connectivity augmentation problems for undirected hypergraphs. After a short introduction on graph edge-connectivity augmentation, we introduce different kinds of objectives that can be prescribed in a hypergraph augmentation problem, and cite the relevant known results. We also show how the covering of certain types of set functions relates to edge-connectivity augmentation. Finally in Section 3.4 we prove a common generalization of the result of Benczúr and Frank [9] and of Fleiner and Jordán [20]. The result (taken from [51]) implies a formula and an algorithm for minimum cardinality k-edge-connectivity augmentation using uniform hyperedges.

3.1.1 Edge-connectivity augmentation of graphs

Initial deep results on the edge-connectivity augmentation of graphs are due to Lovász [55] and to Watanabe and Nakamura [71], on the minimum number of edges needed to be added to a graph to make it k-edge-connected (this is called the *minimum cardinality problem*). Watanabe and Nakamura gave the following characterization:

Theorem 3.1 (Watanabe, Nakamura [71]). *Let $G_0 = (V, E_0)$ be a graph, and $k \geq 2$ an integer. G_0 can be made k-edge-connected by adding at most γ new edges if and only if*

$$\sum_{Z \in \mathcal{F}} (k - d_{G_0}(Z)) \leq 2\gamma \quad \text{for every subpartition } \mathcal{F} \text{ of } V.$$

Not that the theorem does not hold for $k = 1$, but in that case even the minimum cost problem is solvable in polynomial time. For $k \geq 2$ the minimum cost problem is NP-complete. Watanabe and Nakamura showed that a minimum cardinality augmentation can be obtained in polynomial time by repeatedly increasing the edge-connectivity of the graph by one using the minimum number of edges. However, this algorithm is not strongly polynomial.

Frank [29] gave the first strongly polynomial algorithm for this problem. The algorithm relies on the following result concerning degree specified augmentation:

Theorem 3.2. *Let $G_0 = (V, E_0)$ be a graph, $k \geq 2$ an integer, and $m : V \to \mathbb{Z}_+$ a degree specification such that $m(V)$ is even. There is a graph G such that $d_G(v) = m(v)$ for every $v \in V$ and $G_0 + G$ is k-edge-connected if and only if*

$$m(X) \geq k - d_{G_0}(X) \quad \text{for every } \emptyset \neq X \subset V.$$

Furthermore, Frank proved in [29] that the local edge-connectivity augmentation of graphs (i.e. when for every pair $u, v \in V$ there is a local edge-connectivity requirement $r(u, v)$) can also be solved in strongly polynomial time. These proofs use the so-called *splitting-off* technique to solve the degree specified problem. Splitting off at a given node $s \in V$ means deleting edges st_1 and st_2, and adding a new edge t_1t_2. A sequence of splitting-off operations that isolate the node s is called a *complete splitting* at s. The splitting-off operation was originally introduced by Lovász [55] and subsequently developed further by Mader ([58], [59]) and others. Here we cite the result of Mader [58] on splitting-off preserving local edge-connectivity.

Theorem 3.3 (Mader [58]). *Let $G = (V + s, E)$ be a connected graph, where $d(s) \neq 3$ and there is no cut-edge or loop incident to s. Then there are edges $st_1 \in E$ and $st_2 \in E$ such that for the graph $G' = (V, E - \{st_1, st_2\} + \{t_1t_2\})$, $\lambda_{G'}(u, v) = \lambda_G(u, v)$ for every $u, v \in V$.* □

3.1.2 Additional requirements

Many interesting results on edge-connectivity augmentation concern problems where there are additional requirements for the obtained graph besides prescribed connectivity. We cite a few results of this type.

It was proved in [47] that deciding whether there is a complete splitting-off at a specified node of a given simple graph that preserves simplicity is NP-complete. However, Bang-Jensen and Jordán were able to prove the following:

Theorem 3.4 (Bang-Jensen, Jordán [3]). *For any fixed k, there is a polynomial algorithm that computes the minimum number of edges to be added to a given simple graph to make it k-edge-connected, while preserving simplicity.*

Bang-Jensen, Gabow, Jordán, and Szigeti proved that the problem of finding a minimum cardinality augmentation where the new edges are cross-edges of a given partition is solvable in polynomial time:

Theorem 3.5 (Bang-Jensen et al. [1]). *Let $G_0 = (V, E_0)$ be a graph, k a positive integer, and $\mathcal{F}$ a partition of V. An edge is* allowed *if its two endpoints are in different members of $\mathcal{F}$. It is possible to find in polynomial time the minimum number of allowed edges whose addition to G_0 results in a k-edge-connected graph.*

As a special case, given a bipartite graph, one can find the minimum number of edges whose addition results in a k-edge-connected bipartite graph.

Given a planar graph $G = (V+s, E)$, one can ask whether there is a complete splitting-off at s that preserves planarity. Nagamochi and Eades [60] obtained the following result:

Theorem 3.6 (Nagamochi, Eades [60]). *Let $k = 3$ or k be an even integer, and let $G = (V + s, E)$ be a planar graph that is k-edge-connected in V. Then there exists a complete splitting-off at s such that the resulting graph is k-edge-connected in V and planar.*

Recently, Jordán [48] devised a method for solving the simultaneous edge-connectivity augmentation of two graphs by the same edge set:

Theorem 3.7 (Jordán [48]). *Let $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ be graphs, and k, l positive integers. It is possible to find in polynomial time a minimum cardinality edge set E for which the graph $G_1' = (V, E_1 + E)$ is k-edge-connected and the graph $G_2' = (V, E_2 + E)$ is l-edge-connected.*

The method used in the solution of this problem is based on the structural characterization of non-splittable pairs of edges. This characterization also gives a simpler proof of the result of Nagamochi and Eades [60] on planarity-preserving splitting-off.

3.1.3 Hypergraphs: various objectives

Edge-connectivity augmentation problems for hypergraphs have been less intensively studied than their graph counterparts. However, some recent results have shown that they are also worthy of interest. In [11], Cheng gave a formula on the minimum number of graph edges that can be added to an initial $(k-1)$-edge-connected hypergraph so that the resulting hypergraph is k-edge-connected; Bang-Jensen and Jackson [2] extended this result to the case when the initial hypergraph can be arbitrary. More general frameworks, involving the covering of set functions or of families of sets, were proposed by Benczúr and Frank [9] and Fleiner and Jordán [20]. Szigeti [69] solved the local edge-connectivity augmentation problem for hypergraphs when the aim is to minimize the total size of the added hyperedges.

These results show that the objective of minimizing the number of new edges, which we used in the graph case, can be generalized in various directions. In the rest of the chapter we address these possibilities. Of course minimizing the number of new hyperedges is in itself an uninteresting question, since it is always best to use hyperedges containing the whole ground set. One alternative is to consider instead the total size (the sum of the sizes) of the hyperedges; another is to constrain the size of the added hyperedges.

In the following sections we cite the known results about these objectives. We also describe the generalization studied in [9] and [69], that involves covering certain types of set functions by a hypergraph. The chapter is concluded by a new result on covering symmetric crossing supermodular functions by uniform hypergraphs, which extends results in [20] and [9].

3.2 Minimum size augmentation

3.2.1 Local edge-connectivity requirements

Let us first consider the objective of minimizing the sum of the sizes of the added hyperedges, called the *total size*. It turns out that in some sense this problem is actually easier then graph edge-connectivity augmentation. Local edge-connectivity augmentation can be solved, and it does not require tools as sophisticated as Theorem 3.3 of Mader.

Let $r : V^2 \to \mathbb{Z}_+$ be a local edge-connectivity requirement function for which $r(u,v) = r(v,u)$ and $r(v,v) = 0$ for every $u, v \in V$. For a set $X \subseteq V$, let $R(X) := \max\{r(u,v) : u \in X,\ v \in V - X\}$. Szigeti [69] proved the following:

Theorem 3.8 (Szigeti [69]). *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, and r a local edge-connectivity requirement function like above. There is a hypergraph H with hyperedges of* *total size σ such that $\lambda_{H_0+H}(u,v) \geq r(u,v)$ for every $u,v \in V$ if and only if $\sum_{Z\in\mathcal{F}}(R(Z) - d_{H_0}(Z)) \leq \sigma$ holds for every subpartition $\mathcal{F}$ of V.*

The proof also yields a polynomial-time algorithm for the solution.

3.2.2 Covering skew supermodular functions

In the above problem, finding an augmenting hypergraph with minimum total size amounts to finding a hypergraph of minimum total size that covers the set function $R(Z) - d_{H_0}(Z)$. As it was mentioned in Section 2.3, this set function is symmetric and skew supermodular. So a natural generalization of the augmentation problem is to find a hypergraph with minimum total size that covers a given symmetric skew supermodular set function. Actually it was this problem that was solved by Szigeti:

Theorem 3.9 (Szigeti [69]). *Let $p : 2^V \to \mathbb{Z}$ be a symmetric skew-supermodular set function. There is a hypergraph H with hyperedges of total size σ that covers p if and only if $\sigma \geq \sum_{Z\in\mathcal{F}} p(Z)$ for every subpartition $\mathcal{F}$ of V.*

Note that the symmetry of p is not crucial: if p is skew supermodular, then the set function $p'(X) := \max\{p(X), p(V - X)\}$ is symmetric and skew supermodular.

3.3 Augmentation using graph edges

An alternative to minimum total size is to consider minimum cardinality augmentation, but with a constraint on the size of the added hyperedges. The most simple case is when we allow only the addition of graph edges.

3.3.1 NP-completeness of local edge-connectivity augmentation

As we indicated, there is a polynomial algorithm for local edge-connectivity augmentation of hypergraphs with a minimum total size of hyperedges. It was also mentioned that using Theorem 3.3, local edge-connectivity augmentation of graphs using a minimum number of edges is solvable. Though this raised hopes that the minimum cardinality problem for hypergraphs, i.e the minimization of the number of new graph edges added to the initial hypergraph, might also be tractable for local edge-connectivity augmentation, it turned out that this problem is NP-complete. NP-completeness was proved by Cosh, Jackson, and Z. Király for the following special case:

Theorem 3.10 (Cosh et al. [12]). *Let $H_0 = (V, \mathcal{E}_0)$ be a connected hypergraph, $\mathcal{F}$ a partition of V, and γ a non-negative integer. The following problem is NP-complete: decide whether there exists a graph $G = (V, E)$ with γ edges such that $\lambda_{H_0+G}(u, v) \geq 2$ whenever u and v are in the same member of $\mathcal{F}$.*

3.3.2 Global edge-connectivity requirement

When considering k-edge-connectivity augmentation, there is another type of difference in difficulty between graph augmentation and hypergraph augmentation using graph edges. Namely, Watanabe and Nakamura [71] proved that in the graph case there is always a minimum cardinality augmentation which can be obtained by a series of augmentations which optimally increase the edge-connectivity of the graph by 1. However, Benczúr and Cheng have shown that it is not always possible to do this for hypergraphs.

In [11], Cheng gave a formula on the minimum number of graph edges that can be added to an initial $(k-1)$-edge-connected hypergraph such that the resulting hypergraph is k-edge-connected. Bang-Jensen and Jackson [2] extended this result to the case when the initial hypergraph can be arbitrary. Let $c(H)$ denote the number of connected components of the hypergraph H. The min-max theorem is the following:

Theorem 3.11 (Bang-Jensen, Jackson [2]). *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, and k a positive integer. There is a graph G with γ edges such that $H_0 + G$ is k-edge-connected if and only if the following hold:*

$$2\gamma \geq \sum_{Z \in \mathcal{F}} (k - d_{H_0}(Z)) \quad \text{for every sub-partition } \mathcal{F} \text{ of } V, \tag{3.1}$$

$$\gamma \geq c(H_0 - \mathcal{E}'_0) - 1 \quad \text{for every } \mathcal{E}'_0 \subseteq \mathcal{E}_0 \text{ for which } |\mathcal{E}'_0| = k - 1. \tag{3.2}$$

Bang-Jensen and Jackson used a splitting-off technique which is much more complicated then the one of Lovász [55] or Mader [59], but it still gives rise to a polynomial-time algorithm.

3.3.3 Covering symmetric supermodular functions by graphs

For a hypergraph H_0 we can define the set function $p(X) := k - d_{H_0}(X)$ if $\emptyset \neq X \subset V$, and $p(\emptyset) = p(V) := 0$. The k-edge-connectivity augmentation of H_0 by graph edges corresponds to the covering of p by a graph. The set function p is symmetric and crossing supermodular (and $(p)^+$ is positively crossing supermodular). The result of Bang-Jensen and Jackson was generalized in this direction by Benczúr and Frank in [9], where they considered the minimum number of graph edges that can cover a given symmetric, positively crossing

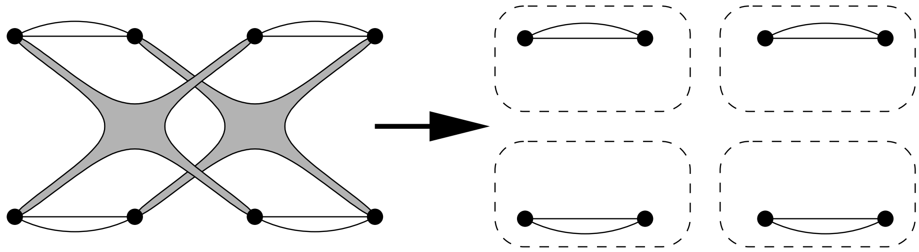

Figure 3.1: A p-full partition with respect to $p = 3 - d_{H_0}$: the components after the deletion of 2 hyperedges

supermodular set function. This more abstract setting provided a better insight into the combinatorial structure underlying the augmentation problem.

If there is a partition $\mathcal{F} = \{V_1, \ldots, V_l\}$ such that $p(\cup_{i\in I} V_i) > 0$ for every $\emptyset \neq I \subset \{1, \ldots, l\}$, then the number of edges of the covering graph is at least $l - 1$, since the edges must connect the members of the partition. Such partitions are called *p-full* and they play an important role in the solution of the problem. For example, when $p(X) = k - d_{H_0}(X)$ if $\emptyset \neq X \subset V$, then the components considered in (3.2) form a p-full partition (as illustrated on Figure 3.1). Benczúr and Frank proved the following theorem:

Theorem 3.12 (Benczúr, Frank [9]). *Let $p : 2^V \to \mathbb{Z}_+$ be a symmetric, positively crossing supermodular set function. There is a graph with γ edges that covers p if and only if the following hold:*

$$2\gamma \geq \sum_{Z\in\mathcal{F}} p(Z) \quad \textit{for every partition } \mathcal{F} \textit{ of } V, \tag{3.3}$$

$$\gamma \geq |\mathcal{F}| - 1 \quad \textit{for every } p\textit{-full partition } \mathcal{F}. \tag{3.4}$$

The core of the proof is an extension of the splitting-off operation to this abstract setting, in order to provide a solution for the corresponding degree-specified problem. Note that in this theorem, in contrast to the skew supermodular case, the requirement of symmetry is important: one cannot make a non-symmetric p symmetric by taking $\max\{p(X), p(V-X)\}$ because this set function is not necessarily positively crossing supermodular.

3.4 Covering by uniform hypergraphs

3.4.1 Covering symmetric crossing families

Another generalization of Cheng's result due to Fleiner and Jordán [20] concerns the covering of symmetric crossing families by uniform hypergraphs. This includes the problem of adding ν-hyperedges (for a fixed ν) to a $(k-1)$-edge-connected hypergraph to make it k-edge-connected. A hypergraph H covers a family $\mathcal{C}$ if $d_H(X) > 0$ for every $X \in \mathcal{C}$. A family $\mathcal{C}$ of sets is *crossing* if for any two crossing sets $X, Y \in \mathcal{C}$, $X \cap Y$ and $X \cup Y$ are in $\mathcal{C}$. Note that if $\mathcal{C}$ is also symmetric, then $X - Y$ and $Y - X$ are in $\mathcal{C}$ as well. If we define $p(X) := 1$ if $X \in \mathcal{C}$ and $p(X) := -\infty$ otherwise, then p is crossing supermodular.

If a hypergraph H_0 is $(k-1)$-edge-connected, then the sets $\emptyset \neq X \subset V$ for which $k - d_{H_0}(X) = 1$ form a symmetric crossing family, and a hypergraph H covers this family if and only if $H_0 + H$ is k-edge-connected.

A partition $\mathcal{F} = \{V_1, \dots, V_l\}$ is called a *full partition* if $\cup_{i\in I} V_i \in \mathcal{C}$ for every $\emptyset \neq I \subset \{1, \dots, l\}$. In other words, it is full if it is p-full for the set function p defined above. Fleiner and Jordán proved the following theorem:

Theorem 3.13 (Fleiner, Jordán [20]). *Let $\mathcal{C}$ be a symmetric crossing family on the ground set V, and $\nu \geq 2$ an integer. Then $\mathcal{C}$ can be covered by γ ν-hyperedges if and only if*

$$\nu\gamma \geq \max\{|\mathcal{F}| : \ \mathcal{F} \subseteq \mathcal{C},\ \mathcal{F} \textit{ is a subpartititon of } V\},$$

and

$$(\nu - 1)\gamma \geq \max\{|\mathcal{F}| - 1 : \ \mathcal{F} \textit{ is a full partititon}\}.$$

Note that the covering by ν-hyperedges is equivalent to the problem of covering by hyperedges of size at most ν, since one can always take hyperedges of maximal size. The proof of Fleiner and Jordán is different from the proofs of the results cited previously, in that it does not solve a degree-specified result using some splitting-off technique, instead it depends on an analysis of the structure of symmetric crossing families.

3.4.2 Covering symmetric supermodular functions by uniform hypergraphs

The main result of this chapter is a common generalization of the above mentioned results in [9] and [20] (Theorems 3.12 and 3.13), based on the approach of Benczúr and Frank. We give a min-max formula on the minimum number of ν-hyperedges that can cover a given symmetric, positively crossing supermodular set function. As in [9], the substantial part of the proof is a solution of the degree-specified problem (i.e. when the degree of each node $v \in V$ is a prescribed value $m(v)$), which then easily leads to a min-max formula on the minimum number of new hyperedges needed.

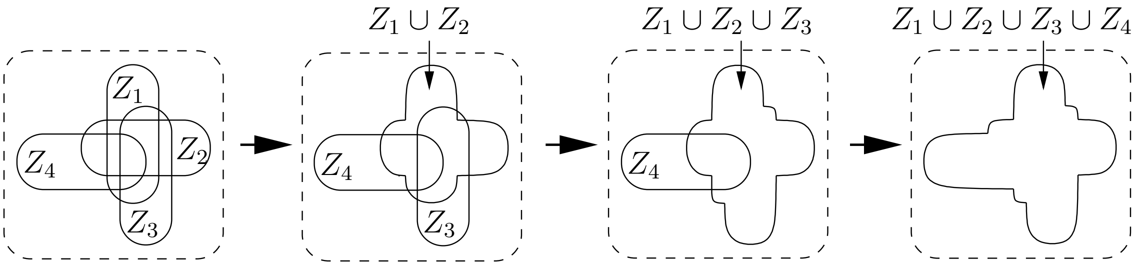

Figure 3.2: Repeated uncrossing

Let $p : 2^V \to \mathbb{Z}_+$ be a symmetric, positively crossing supermodular set function. It was remarked in Chapter 2 that in this case inequality (2.9) holds for every pair (X, Y) for which $p(X) > 0$, $p(Y) > 0$, and $X - Y$ and $Y - X$ are non-empty.

Another tool related to supermodularity that we will use is repeated uncrossing (Figure 3.2). Let $Z_1, \dots, Z_l$ be subsets of V such that $p(Z_i) > 0$ $(i = 1, \dots, l)$, $\cup_{i=1}^l Z_i \neq V$, $Z_{j+1} \cap (\cup_{i=1}^j Z_i) \neq \emptyset$, and $p(Z_{j+1} \cap (\cup_{i=1}^j Z_i)) \leq 1$ $(j = 1, \dots, l-1)$. Then by using (2.2) we get

$$p(\cup_{i=1}^l Z_i) \geq p(\cup_{i=1}^{l-1} Z_i) + p(Z_l) - 1, \tag{3.5}$$

and by induction

$$p(\cup_{i=1}^l Z_i) \geq p(\cup_{i=1}^j Z_i) + \sum_{i=j+1}^{l} p(Z_i) - (l - j), \tag{3.6}$$

$$p(\cup_{i=1}^l Z_i) \geq \sum_{i=1}^{l} p(Z_i) - (l - 1). \tag{3.7}$$

Let $\nu \geq 2$ be an integer, and $m : V \to \mathbb{Z}_+$ a degree specification such that ν divides $m(V)$. First we consider the problem of finding a ν-regular hypergraph satisfying this degree-specification that covers the set function p. There is an obvious lower bound on $m(V)$ that was implied implicitly by the conditions of Theorems 3.12 and 3.13, but needs to be stated explicitly here: the number of new hyperedges must be at least $\max_{X \subseteq V} p(X)$.

As in Theorem 3.12, there will be a condition featuring p-full partitions. We call a partition $\mathcal{F} = \{V_1, \dots, V_l\}$ *p-full* if $l > \nu$ and

$$p(\cup_{i \in I} V_i) > 0 \quad \text{for every } \emptyset \neq I \subset \{1, \dots, l\}. \tag{3.8}$$

We always assume that the partition members are indexed so that $m(V_1) \leq m(V_2) \leq \dots \leq m(V_l)$. Suppose that a ν-uniform hypergraph covers p. If we contract the sets $V_1, \dots, V_l$, then the contracted hypergraph (which is still ν-uniform since multiplicities are taken into account) must be connected, therefore it needs to have at least $\frac{l-1}{\nu-1}$ hyperedges. A p-full partition is called a *deficient partition* if

$$\frac{l-1}{\nu-1} > \frac{m(V)}{\nu}.$$

Theorem 3.14. *Let $p : 2^V \to \mathbb{Z}_+$ be a symmetric, positively crossing supermodular set function, $\nu \geq 2$ an integer, and $m : V \to \mathbb{Z}_+$ a degree specification such that $\nu \mid m(V)$. There is a ν-uniform hypergraph H covering p such that $d_H(v) = m(v)$ for every $v \in V$ if and only if the following are true:*

$$m(X) \geq p(X) \quad \text{for every } X \subseteq V, \tag{3.9}$$

$$\frac{m(V)}{\nu} \geq p(X) \quad \text{for every } X \subseteq V, \tag{3.10}$$

$$\text{There are no deficient partitions.} \tag{3.11}$$

Proof. The necessity of the conditions is easily verifiable. We prove sufficiency using induction on $|V| + m(V)$. If $m(V) = \nu$, conditions (3.9) and (3.10) are clearly sufficient, so we can assume that $m(V) \geq 2\nu$. A ν-uniform hypergraph is called *feasible* if it matches the degree specification and covers p. First we show that if there is a set $X \subseteq V$ such that $m(X) = p(X) = 1$ and $|X| > 1$, then there exists a feasible hypergraph. The contraction of X leads to a modified problem: $V' := V - X + v_X$, $m'(v) := m(v)$ if $v \in V' - v_X$, $m'(v_X) := 1$, $p'(Y) := p(Y)$ if $v_X \notin Y$, and $p'(Y) := p((Y - v_X) \cup X)$ if $v_X \in Y$. Conditions (3.9)–(3.11) are satisfied by m' and p', and p' is symmetric and positively crossing supermodular, so by induction there is a ν-uniform hypergraph $H' = (V', \mathcal{E}')$ with degree vector m', that covers p'. This hypergraph naturally defines a ν-uniform hypergraph $H = (V, \mathcal{E})$ with degree vector m; we claim that H covers p.

Suppose that some $Y \subseteq V$ is deficient, i.e. $d_H(Y) < p(Y)$. Then $p(Y) \geq 1$ and $m(Y) \geq 2$, so $Y \not\subseteq X$. Furthermore, Y must separate X, otherwise there would be a corresponding deficient set in the contracted problem. If $m(X \cap Y) > 0$, then we may assume that $V - (X \cup Y) \neq \emptyset$, because otherwise $0 = m(X - Y) < p(Y) = p(X - Y)$. Using (2.2), $p(X \cup Y) \geq p(X) + p(Y) - p(X \cap Y) \geq p(Y) > d_H(Y) = d_H(X \cup Y)$, so $X \cup Y$ would be deficient. If $m(X \cap Y) = 0$, then $p(Y - X) \geq p(X) + p(Y) - p(X - Y) \geq p(Y) > d_H(Y) = d_H(Y - X)$ according to (2.9), so $Y - X$ would be deficient.

From now on it is assumed that if $m(X) = p(X) = 1$ for some $X \subseteq V$, then $|X| = 1$; these singletons are called *special singletons*. The set of special singletons is denoted by S, and it is considered as a subset of V.

We define an operation called *splitting-off*, which is an analogue of the splitting-off operation for graphs. A ν-hyperedge e can be split off from (p, m) if $|e \cap v| \leq m(v)$ for every $v \in V$. For such a hyperedge let

$$m^e(v) := m(v) - |e \cap v|, \tag{3.12}$$

$$p^e(X) := \begin{cases} (p(X)-1)^+ & \text{if } e \text{ enters } X, \\ p(X) & \text{otherwise.} \end{cases} \tag{3.13}$$

We say that (m^e, p^e) is obtained from (m, p) by *splitting off* the hyperedge e. A splitting-off operation is *feasible* if (3.9), (3.10), and (3.11) are true for m^e and p^e. It is easy to see that p^e is symmetric and positively crossing supermodular; so after the execution of a feasible splitting-off, by induction there exists a ν-uniform hypergraph H' with degree vector m^e that covers p^e . By adding the hyperedge e to H' we obtain a feasible hypergraph H.

The rest of the proof consists of showing that a feasible splitting-off always exists. We define the following families of sets:

$$\mathcal{B}_1 := \{\emptyset \neq X \subseteq V : \ m(X) - p(X) \leq \nu - 2, \ p(Y) < p(X) \ \forall Y \subset X\},$$

$$\mathcal{B}_2 := \{X \subseteq V : \ m(X) - p(X) = \nu - 1, \ p(X) > 0, \ p(Y) \leq p(X) \ \forall Y \subset X\},$$

$$\mathcal{B}_3 := \{X \subseteq V : \ p(X) = \frac{m(V)}{\nu}, \ p(Y) < \frac{m(V)}{\nu} \ \forall Y \subset X\}.$$

Let e be a ν-hyperedge that can be split off, and suppose that m^e and p^e violates (3.9) because there is a set X entered by e such that $p(X) - 1 > m(X) - |e \cap X|$. Notice that $m(X') - |e \cap X'| \leq m(X) - |e \cap X|$ for every $X' \subseteq X$. Thus there must be a set $X' \in \mathcal{B}_1$ such that $X' \subseteq X$ and $p(X') - 1 > m(X') - |e \cap X'|$. Similarly, if m^e and p^e violates (3.9) because there is a set X such that $e \subseteq X$ and $p(X) = m(X) - (\nu - 1)$, then either there is a set $X' \in \mathcal{B}_2$ such that $X' \subseteq X$ and $e \subseteq X'$, or there is a set $X' \in \mathcal{B}_1$ such that $X' \subseteq X$ and $p(X') - 1 > m(X') - |e \cap X'|$.

A similar argument shows that if m^e and p^e violate (3.9) then there is a set $X \in \mathcal{B}_3$ such that $e \subseteq V - X$. We obtained the following:

Claim 3.15. *The inequalities (3.9) and (3.10) hold for m^e and p^e if and only if*

$$|e \cap X| \leq m(X) - p(X) + 1 \quad \textit{for every } X \in \mathcal{B}_1, \tag{3.14}$$

$$|e \cap X| \leq \nu - 1 \quad \textit{for every } X \in \mathcal{B}_2, \tag{3.15}$$

$$|e \cap X| \geq 1 \quad \textit{for every } X \in \mathcal{B}_3. \tag{3.16}$$

Furthermore, if $m^e(X) < p^e(X)$ for some X, then there is a subset X' of X such that either $X' \in \mathcal{B}_1$ and it violates (3.14), or $X' \in \mathcal{B}_2$ and it violates (3.15). □

In order to formulate necessary and sufficient conditions for m^e and p^e to satisfy (3.11), we call a p-full partition $\{V_1, \dots, V_l\}$ *critical* if

$$\frac{l-1}{\nu-1} > \frac{m(V)}{\nu} - 1.$$

For a partition $\mathcal{F}$, let $s(\mathcal{F})$ denote the number of special singleton members of $\mathcal{F}$. A critical partition $\mathcal{F}$ is called *proper* if $s(\mathcal{F}) \geq 3$. Critical partitions have the following properties:

Claim 3.16. *If $\mathcal{F} = \{V_1, \ldots, V_l\}$ is a critical partition, then $2l - 2 \geq m(V)$, thus $s(\mathcal{F}) \geq m(V_l)$. In particular, $s(\mathcal{F}) \geq 2$ for every critical partition, and the partition is proper if $m(V_l) \geq 3$.*

Proof. The partition is critical and $\frac{m(V)}{\nu}$ is an integer, so

$$m(V) \leq \nu \left(\frac{l-2}{\nu - 1} + 1 \right) = 2\nu + \frac{\nu}{\nu - 1}(l - \nu - 1) \leq$$
$$\leq 2\nu + 2(l - \nu - 1) = 2l - 2,$$

since p-fullness implies that $l > \nu$. □

Claim 3.17. *A partition $\{V_1, \ldots, V_l\}$ is critical if and only if $l > \nu$, $\frac{l-1}{\nu-1} > \frac{m(V)}{\nu} - 1$, $p(V_1) = 1$, and $p(V_1 \cup V_i) \geq 1$ $(i = 2, \ldots, l)$. If the partition is critical and U is the union of some partition members such that $V_1 \subseteq U$ and $V_2 \cap U = \emptyset$, then $p(U) = 1$.*

Proof. Let $\{V_1, \ldots, V_l\}$ be a partition with the above properties, and let $U = \cup_{i \in I} V_i$ where $1 \in I$ and $|I| \leq l - 1$. We can use inequality (3.7) for the sets $\{V_1 \cup V_i : i \in I\}$ to show that $p(U) > 0$. The symmetry of p implies that $p(V - U) > 0$ for all such U, so the partition is p-full. If $V_1 \subseteq U$ and $V_2 \cap U = \emptyset$, then (3.6) gives $1 = p(V_1) \leq p(U) \leq p(V - V_2) = 1$. □

Claim 3.18. *Let e be a ν-hyperedge which satisfies (3.14)–(3.16). Then m^e and p^e satisfy (3.11) if and only if*

$$e \not\subseteq X \text{ for any member } X \text{ of any proper critical partition.} \tag{3.17}$$

Proof. If $e \subseteq X$ for a member X of a critical partition, then the partition remains p-full after the splitting off of e, hence the partition becomes deficient. To see the converse, observe that only critical partitions can become deficient partitions after the splitting-off. If $e \not\subseteq X$ for every member X of a given critical partition $\mathcal{F}$, then it is easy to see that there exists a set U entered by e which is the union of some members of $\mathcal{F}$ including exactly 1 special singleton member. According to Claim 3.17, $p(U) = 1$, so $p^e(U) = 0$ and $\mathcal{F}$ is not p-full after the splitting-off. It remains to show that if $e \subseteq X$ for a member X of a non-proper critical partition, then some set violates (3.14) or (3.15). But $m(X) \leq 2$ according to Claim 3.16, so $0 = m^e(X) < p(X) = p^e(X)$ which by Claim 3.15 implies that (3.14) or (3.15) is violated for some subset of X. □

It suffices to show the existence of a ν-hyperedge e for which $\chi_e \leq m$ and which satisfies (3.14), (3.15), (3.16), and (3.17). First we consider only (3.14) and (3.16):

$$Q := \{e \in \mathbb{Z}_+^V : \; \chi_e \leq m; \; |e| = \nu; \\ |e \cap X| \leq m(X) - p(X) + 1 \; \forall X \in \mathcal{B}_1; \; |e \cap X| \geq 1 \; \forall X \in \mathcal{B}_3\}.$$

Observe that $p(X) > 0$ for every $X \in \mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3$, so inequalities (2.2) and (2.9) can be used for these sets.

Claim 3.19. *The family $\mathcal{B}_1 \cup \mathcal{B}_3$ is laminar. The sets in $\mathcal{B}_3$ are pairwise disjoint, and if $X \in \mathcal{B}_1$ and $Y \in \mathcal{B}_3$ are not disjoint, then $X \subseteq Y$.*

Proof. If $X, Y \in \mathcal{B}_1 \cup \mathcal{B}_3$, and $X - Y, Y - X, X \cap Y \neq \emptyset$, then $p(X) \leq p(X - Y)$ or $p(Y) \leq p(Y - X)$ by (2.9), which contradicts the definition of $\mathcal{B}_1$ and $\mathcal{B}_3$. If $X \in \mathcal{B}_1$ and $Y \in \mathcal{B}_3$, then $p(Y) \geq p(X)$, so $Y \not\subset X$ according to the definition of $\mathcal{B}_1$. □

Claim 3.20. *Q is non-empty.*

Proof. We use the following algorithm to find an element e of Q:

1. Construct a hyperedge e' by choosing one node from every $X \in \mathcal{B}_3$.
2. Fix an arbitrary ordering of the nodes of V, and consider them one by one. Increase the multiplicity in e' of each node to the maximum value with which the obtained hyperedge does not violate conditions of type (3.14) and its size is at most ν.

It follows from Claim 3.19 that this algorithm finds a hyperedge $e \in Q$ if and only if the following hold:

(i) $|\mathcal{B}_3| \leq \nu$,

(ii) $m(V - \cup_{i=1}^t X_i) + \sum_{i=1}^t (m(X_i) - p(X_i) + 1) \geq \nu$ for every sub-partition $\{X_1, \ldots, X_t\}$.

The first condition holds since $m(X) \geq \frac{m(V)}{\nu}$ for every $X \in \mathcal{B}_3$, and they are disjoint by Claim 3.19. The second condition is clearly true if $t \geq \nu$. If $t < \nu$, then $m(V - \cup X_i) + \sum(m(X_i) - p(X_i) + 1) \geq m(V - \cup X_i) + \sum(m(X_i) - \frac{m(V)}{\nu} + 1) = (\nu - t)\frac{m(V)}{\nu} + t \geq \nu$, the last inequality being valid because $m(V) \geq \nu$. □

Obviously, if a ν-hyperedge e can be feasibly split off, then it is in Q. The converse is generally not true; however, it turns out to be true when $\mathcal{B}_3 \neq \emptyset$:

Lemma 3.21. *If $\mathcal{B}_3 \neq \emptyset$, then any ν-hyperedge $e \in Q$ can be feasibly split off.*

Proof. Let $X \in \mathcal{B}_3$ and $e \in Q$. First suppose that there is a set $Y \in \mathcal{B}_2$ such that $e \subseteq Y$. Then $X \cap Y \neq \emptyset$, and one of $X - Y$ and $Y - X$ is empty, otherwise (2.9) would imply that either $p(X - Y) \geq p(X)$ or $p(Y - X) > p(Y)$, contrary to the definition of $\mathcal{B}_3$ and $\mathcal{B}_2$. If $Y \subseteq X$, then there is a set $X' \in \mathcal{B}_3$ such that $X' \subseteq V - X \subseteq V - Y$ because of the symmetry of p; if $X \subseteq Y$, then $Y \in \mathcal{B}_2$ implies that $p(V - Y) = p(Y) \geq p(X) = \frac{m(V)}{r}$, so $\mathcal{B}_3$ would again contain a set $X' \subseteq V - Y$. This is impossible since $e \in Q$, which requires $|e \cap X'| \geq 1$.

Now suppose that a proper critical partition $\mathcal{F} = \{V_1, \ldots, V_l\}$ becomes deficient after the splitting-off. Since $\mathcal{B}_3$ contains at least two disjoint sets, it contains a set Z that is disjoint from at least two special singleton members of $\mathcal{F}$, say V_1 and V_2. First we show by using inequalities (2.2) and (2.9) that a member of $\mathcal{F}$ cannot separate Z. If V_i separates Z, then $p(V_1 \cup V_i - Z) \geq p(V_1 \cup V_i) + p(Z) - p(Z - V_i) > p(V_1 \cup V_i) = 1$. Let $U := V - V_i - V_2$; then $p(U \cup Z) \geq p(U) + p(Z) - p(Z \cap U) > p(U) = 1$. Thus $1 = p(V_i \cup U) \geq p(V_1 \cup V_i - Z) + p(U \cup Z) - p(V_1) > 1$, a contradiction.

We can conclude that there is a partition member V_i such that $Z \subseteq V_i$. This implies $m^e(V_i) \geq m^e(Z) \geq p^e(Z) = \frac{m^e(V)}{\nu}$, so $m^e(V) \geq l - 1 + \frac{m^e(V)}{\nu}$. But then $\frac{l-1}{\nu-1} \leq \frac{m^e(V)}{\nu}$, so $\mathcal{F}$ could not become deficient after the splitting-off. □

By Claim 3.20 and Lemma 3.21 we may assume that $\mathcal{B}_3 = \emptyset$. To handle condition (3.17), we will use some information about the structure of proper critical partitions. This information is based on an auxiliary graph G defined on the special singletons: $G = (S, F)$, where $uv \in F$ if and only if $p(\{u, v\}) = 1$. Note that by Claim 3.17 the special singleton members of a proper critical partition form a clique in G.

Claim 3.22. *If $X \in \mathcal{B}_1$ and $|X| \geq 2$, then $d_G(X) = 0$.*

Proof. Suppose that uv is an edge of G such that $u \in X$, $v \notin X$. Then by (2.9), $p(X - u) \geq p(X) + p(\{u, v\}) - p(v) = p(X)$, which contradicts $X \in \mathcal{B}_1$. □

Claim 3.23. *Every proper critical partition $\mathcal{F} = \{V_1, \ldots, V_l\}$ can be refined by separating some special singletons in such a way that the resulting proper critical partition $\mathcal{F}_S$ (the* S-refinement *of $\mathcal{F}$) has the following properties:*

- *If $d_G(X) > 0$ for a member $X \in \mathcal{F}_S$, then X is a special singleton,*
- *The set of special singleton members of $\mathcal{F}_S$ defines a component of G that is a clique (which we call the* S-clique *of $\mathcal{F}$).*

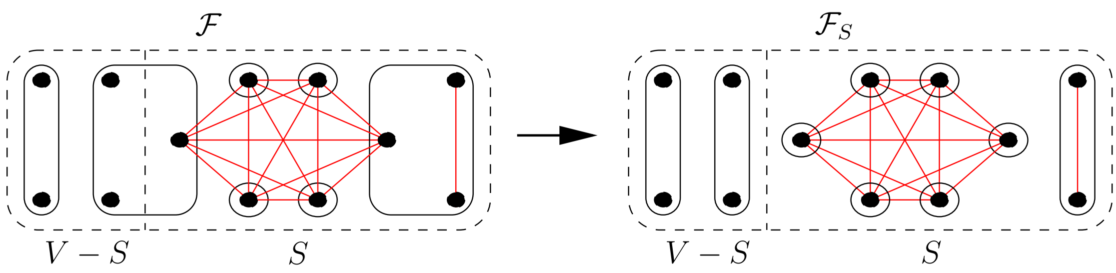

Figure 3.3: The S-refinement of $\mathcal{F}$ (and the auxiliary graph G)

Proof. We prove that if there is an edge uv of G such that $u \in V_i$, $v \notin V_i$ for some non-singleton partition member V_i, then a critical partition is obtained if we replace V_i in $\mathcal{F}$ by $\{u\}$ and $V_i - u$. $\mathcal{F}$ is proper, so we can assume that $v \neq V_1$ and $v \neq V_2$. According to Claim 3.17, we have to show that $p(V_1 \cup V_i - u) > 0$ and $p(V_1 + u) > 0$.

By Claim 3.17, $p(V_1 \cup V_i) = 1$, so $p(V_1 \cup V_i - u) \geq p(V_1 \cup V_i) + p(\{u, v\}) - p(\{v\}) = 1$. Similarly, $p(V_2 \cup V_i - u) \geq 1$; thus $p(V_1 + u) \geq p(V_1 \cup V_i) + p(V_2 \cup V_i - u) - p(V_2) \geq 1$. As a consequence, the replacement of V_i by $\{u\}$ and $V_i - u$ results in a proper critical partition.

By repeating this step as many times as possible, we obtain a proper critical partition $\mathcal{F}_S$ with the required properties. □

Among the components of G that are cliques, let $K^* \subseteq S$ be one of maximal size. The idea behind the next steps of the proof is that a member of a critical partition cannot contain too many nodes of K^* because then the S-refinement would be deficient; therefore we try to choose our splittable hyperedge e so that it would contain many nodes of K^*. More precisely, we first choose a $(\nu - 1)$-hyperedge e' using the following algorithm:

1. Fix an ordering $v_1, \dots, v_n$ of the nodes of V such that the nodes of K^* come first.

2. Consider the nodes one by one in the order $v_1, \dots, v_n$. Let the multiplicity of v_i in e' be the maximum value with which the hyperedge defined so far on $\{v_1, \dots, v_i\}$ does not violate conditions of type (3.14) and its size is at most $\nu - 1$.

Let $W := \{v \in V : m(v) > |e' \cap v|\}$. Analogously to the proof of Claim 3.20, it is easy to see that the size of e' is $\nu - 1$, and there is node $w \in W$ such that $e' + w \in Q$. We will show that actually there exists a node $w \in W$ such that the ν-hyperedge $e = e' + w$ can be feasibly split off.

Lemma 3.24. *If $w \in W$, and the splitting off of the hyperedge $e = e' + w$ results in a deficient partition, then either inequality (3.14) or inequality (3.15) is violated.*

Proof. We can assume that $|K^*| \geq 3$, otherwise there are no proper critical partitions. Claim 3.22 implies that K^* is either disjoint from or subset of any non-singleton set in $\mathcal{B}_1$, so the construction of e' gives that $|e' \cap K^*|$ is the minimum of the following three values:

$$|K^*|, \tag{3.18}$$

$$\nu - 1, \tag{3.19}$$

$$\min_{K^* \subseteq X \in \mathcal{B}_1} (m(X) - p(X) + 1). \tag{3.20}$$

Suppose that a proper critical partition $\mathcal{F}$ becomes deficient after the splitting off of e. Let V_i be the member of $\mathcal{F}$ for which $e \subseteq V_i$; then $|K^* \cap V_i| \geq |e' \cap K^*| > 0$. If $K^* \subseteq V_i$ and $m(V_i - K^*) > 0$, then $s(\mathcal{F}) < m(V_i)$, given that by Claim 3.23 the special singleton members are in the same component of G whose size is at most $|K^*|$. This contradicts the criticality of $\mathcal{F}$ according to Claim 3.16.

If $K^* \subseteq V_i$ and $m(V_i - K^*) = 0$, then the set $Z := \{v \in V_i : \ v \in e\}$ is a subset of K^*, hence $p(\{u, v\}) = 1$ for every $u, v \in Z$. Let $v \in Z$; by using (3.7) on the sets $\{\{v, u\} : \ u \in Z - v\}$ we get that $p(Z) > 0$. But then $m^e(Z) < p^e(Z)$, so by Claim 3.15 a subset of Z violates (3.14) or (3.15).

Now suppose that $K^* \not\subseteq V_i$, and let $\mathcal{F}_S$ denote the S-refinement of $\mathcal{F}$. By Claim 3.23, K^* is the set of special singletons that are members of $\mathcal{F}_S$ (and thus the special singleton members of $\mathcal{F}$ are in K^*). If $e' \subseteq K^*$, then there are at least $\nu - 1$ nodes in V_i that are special singleton members of $\mathcal{F}_S$, so $\frac{|\mathcal{F}_S|-1}{\nu-1} \geq \frac{|\mathcal{F}|+(\nu-1)-1}{\nu-1} \geq \frac{|\mathcal{F}|-1}{\nu-1} + 1 > \frac{m(V)}{\nu}$ by the criticality of $\mathcal{F}$, which means that $\mathcal{F}_S$ would violate (3.11).

If $|e' \cap K^*|$ is determined by (3.20), then there is a set $X \in \mathcal{B}_1$ such that $K^* \subseteq X$, $V_i \not\subseteq X$, and $|K^* \cap V_i| \geq m(X) - p(X) + 1$. Let U denote the union of the members of $\mathcal{F} - \{V_i\}$ that are not special singletons. Then $s(\mathcal{F}) = |K^* - U| - |K^* \cap V_i| \leq |K^* - U| - m(X) + p(X) - 1 < p(X) - m(X \cap U)$. By Claim 3.16, this would imply that $m(V_i) < p(X) - m(X \cap U)$; we show that this is impossible. If $U \subseteq X$, then $m(V_i) \geq p(V_i - X) = p(V - X) = p(X) \geq p(X) - m(X \cap U)$. If $U \not\subseteq X$, then by using (2.2) on X and $V_1 \cup U$ we get that $m(V_i) \geq p(V_i - X) = p(U \cup X) \geq p(X) + p(V_1 \cup U) - p(X \cap (V_1 \cup U)) \geq p(X) + 1 - m(X \cap (V_1 \cup U)) = p(X) - m(X \cap U)$. □

Lemma 3.24 implies that condition (3.17) can be ignored when choosing an appropriate node $w \in W$. We have already seen that $e = e' + w \in Q$ is a necessary condition for the feasibility of the splitting-off. In addition to that, $e \not\subseteq X$ must hold for every set $X \in \mathcal{B}_2$; however, if $e' \not\subseteq X$ or there is a set $Y \in \mathcal{B}_1$ such that $X \subseteq Y$ and $m(Y) - p(Y) = \nu - 2$, then this automatically follows from $e = e' + w \in Q$. So we call a set $X \in \mathcal{B}_2$ *critical* if $e' \subseteq X$ and there is no $Y \in \mathcal{B}_1$ such that $X \subseteq Y$ and $m(Y) - p(Y) = \nu - 2$. A ν-hyperedge

$e = e' + w$ can be feasibly split off if and only if $e \in Q$ and there is no critical set containing w. Since e' was chosen so that $e' + w \in Q$ for at least one $w \in W$, we can assume that there is at least one critical set in $\mathcal{B}_2$.

Claim 3.25. *Let $X \in \mathcal{B}_2$ be a critical set. If $Y \in \mathcal{B}_1$, then one of $X \cap Y$, $X - Y$, and $Y - X$ is empty. If $w \in W - X$, then $e' + w \in Q$.*

Proof. If $X \in \mathcal{B}_2$, $Y \in \mathcal{B}_1$, and $X \cap Y, X - Y, Y - X$ are non-empty, then $p(X - Y) \leq p(X)$ and $p(Y - X) < p(Y)$ by the definition of $\mathcal{B}_1$ and $\mathcal{B}_2$, which contradicts (2.9). If $w \in W - X$, then $e' + w \in Q$ unless there is a set $Y \in \mathcal{B}_1$ such that $w \in Y$ and $|e' \cap Y| = m(Y) - p(Y) + 1$. But then $X \cap Y \neq \emptyset$, so $X \subseteq Y$, which contradicts the criticality of X. □

Suppose indirectly that every $w \in W$ is in a critical set. Consider a family $\mathcal{Z} = \{Z_1, \dots, Z_t\}$ of maximal critical sets, such that every $w \in W$ is in at least one of them, and the family has minimal number of members. The following sequence of claims demonstrates that the existence of such a family leads to a contradiction. Let $Z := \cap_{i=1}^t Z_i$.

Claim 3.26. *If $i \neq j$, then $m(Z_i \cap Z_j) = m(Z) = \nu - 1$, and $p(Z_i) = m(Z_i - Z) = m(Z_i - Z_j) = p(Z_i - Z_j)$.*

Proof. We know that $m(Z_i \cap Z_j) \geq m(Z) \geq |e' \cap Z| = \nu - 1$. Since $Z_i \in \mathcal{B}_2$, we have $p(Z_i) = m(Z_i) - (\nu - 1)$, so (2.9) gives that $0 \leq p(Z_i - Z_j) + p(Z_j - Z_i) - p(Z_i) - p(Z_j) \leq m(Z_i - Z_j) + m(Z_j - Z_i) - p(Z_i) - p(Z_j) = 2(\nu - 1) - 2m(Z_i \cap Z_j)$. This is possible only if equality holds throughout, so $m(Z_i \cap Z_j) = \nu - 1$, and $m(Z_i - Z_j) = p(Z_i - Z_j)$. □

Clearly $|\mathcal{Z}| \geq 2$ and $\frac{m(V)}{\nu} \geq 2$, since $m(V - Z_i) > 0$ for every i. Suppose that $|\mathcal{Z}| = 2$. Then $p(Z_1) = m(Z_1 - Z_2) = m(V - Z_2) \geq p(Z_2)$ and vice versa, so $p(Z_1) = p(Z_2) = m(Z_1 - Z_2) = m(Z_2 - Z_1) = \frac{m(V) - (\nu - 1)}{2}$. This value can be integer only if $\nu \geq 3$, but then $\frac{m(V)}{\nu} \leq \frac{m(V) - (\nu - 1)}{2} = p(Z_1)$, contradicting $\mathcal{B}_3 = \emptyset$. Therefore we may assume that $|\mathcal{Z}| \geq 3$.

Claim 3.27. *$\cup_{i=1}^t Z_i = V$, $Z_i - Z$ is a special singleton for every i, and $p(Z_i \cup Z_j - Z) > 0$ for every i, j.*

Proof. For a set of indices $I \subseteq \{1, \dots, t\}$, let $U_I := \cup_{i \in I} Z_i$. We have seen that $p(Z_i) = m(Z_i - Z)$ for every i. If X is a subset of V for which $m(X) = |e' \cap X| = \nu - 1$, then $p(X) \leq m(X) - |e' \cap X| + 1 = 1$. In particular, this holds for every $Z \subseteq X \subseteq \cup_{i \neq j}(Z_i \cap Z_j)$. Thus we get from (3.7) that $p(U_I) \geq m(U_I - Z) - (|I| - 1) > 0$ if $U_I \neq V$. For $I := \{1, \dots, t\}$ this implies that $\cup_{i=1}^t Z_i = V$, since $p(\cup_{i=1}^t Z_i) \leq m(V - \cup_{i=1}^t Z_i) = 0$. Now let $I := \{1, \dots, t\} - \{i_0\}$, where i_0 is chosen so that $m(Z_{i_0} - Z)$ is minimal. Then $m(Z_{i_0} - Z) \geq p(U_I) \geq m(U_I - Z) - (t - 2)$. Since $|\mathcal{Z}| \geq 3$, this is only possible if $m(Z_i - Z) = 1$ for every i. Thus $Z_i - Z_j$ is a special singleton for every $i \neq j$, which implies that $Z_i - Z$ is a special singleton for every i. Finally, if $i \neq j$, then by setting $I := \{1, \dots, t\} - \{i, j\}$, we obtain $p(Z_i \cup Z_j - Z) = p(U_I) > 0$. □

Claim 3.28. *$\mathcal{F} := \{Z, Z_1 - Z, \dots, Z_t - Z\}$ is a proper critical partition.*

Proof. The partition has size $l = m(V) - (\nu - 1) + 1 \geq \nu + 2$ since $m(V) \geq 2\nu$, so $s(\mathcal{F}) \geq \nu + 1$, and $\frac{l-1}{\nu-1} = \frac{m(V)-(\nu-1)}{\nu-1} > \frac{m(V)}{\nu} - 1$. If X is the union of two partition members, then Claim 3.27 implies that $p(X) > 0$; therefore $\mathcal{F}$ is a proper critical partition by Claim 3.17. □

$\mathcal{F}$ has at least $\nu+1$ special singleton members. According to Claim 3.23, these all belong to the same component of G that is a clique, therefore $|K^*| \geq \nu + 1$. This means that $K^* \nsubseteq Z$, so K^* must be the S-clique of $\mathcal{F}$. The value of $|e' \cap K^*|$ is not determined by (3.20): if $K^* \subseteq X \in \mathcal{B}_1$, then $Z_i \subseteq X$ for every i by Claim 3.25, which is not possible. It follows that $e' \subseteq K^*$.

To prove that the existence of $\mathcal{Z}$ contradicts (3.11), we consider the S-refinement $\mathcal{F}_S$ of $\mathcal{F}$. The properties of S-refinements stated in Claim 3.23 imply that every member of $\mathcal{F}_S$ is a special singleton. However, such a p-full partition would be a deficient partition, since $\frac{m(V)-1}{\nu-1} > \frac{m(V)}{\nu}$ if $m(V) \geq 2\nu$.

We proved that there is a node w such that the ν-hyperedge $e = e' + w$ can be feasibly split off. This concludes the proof of Theorem 3.14. □

3.4.3 Minimum cardinality augmentation

As it is the case with many edge-connectivity augmentation results, the characterization of the degree-specified problem in Theorem 3.14 can be used in a straightforward way to prove a min-max theorem on the corresponding minimum cardinality problem. Recall that a partition $\{V_1, \dots, V_l\}$ is called *p-full* if $l > \nu$ and $p(\cup_{i\in I} V_i) > 0$ for every $\emptyset \neq I \subset \{1, \dots, l\}$.

Theorem 3.29. *Let $p : 2^V \to \mathbb{Z}_+$ be a symmetric, positively crossing supermodular set function, and $\nu \geq 2$ an integer. There is a ν-uniform hypergraph with γ hyperedges that covers p if and only if the following hold:*

$$\nu\gamma \geq \sum_{i=1}^{t} p(X_i) \quad \textit{for every partition } \{X_1, \dots, X_t\}, \tag{3.21}$$

$$\gamma \geq p(X) \quad \textit{for every } X \subseteq V, \tag{3.22}$$

$$\gamma \geq \frac{l-1}{\nu-1} \quad \textit{if there is a p-full partition with } l \textit{ members.} \tag{3.23}$$

Proof. The conditions are clearly necessary for the existence of a ν-uniform hypergraph that covers p. We prove sufficiency for a fixed γ. Let $m' : V \to \mathbb{Z}_+$ be a vector that satisfies (3.9) such that $m'(V)$ is minimal, and let $W := \{v \in V : m'(v) > 0\}$. Then for every node $v \in W$, there exists a set X for which $v \in X$ and $m'(X) = p(X)$; sets with the latter property are called *tight*. There is a family of tight sets covering every node $v \in W$; let $\mathcal{F}$ be such a family with $|\mathcal{F}|$ minimal. If $X \in \mathcal{F}$ and $Y \in \mathcal{F}$ are not disjoint, then $X \cup Y = V$, otherwise $X \cup Y$ would be tight according to (2.2), which would contradict the minimality of $|\mathcal{F}|$, as X and Y could be replaced in $\mathcal{F}$ by $X \cup Y$. The symmetry of p implies that $X - Y = V - Y$ and $Y - X = V - X$ are both tight and $m'(X \cap Y) = 0$, so $X - Y$ and $Y - X$ cover every node of W. We can conclude that there is always a partition $\{X_1, \dots, X_t\}$ for which $\sum_{i=1}^{t} p(X_i) = m'(V)$. It follows from (3.21) that $r\gamma \geq m'(V)$.

We can obtain a degree specification $m : V \to \mathbb{Z}_+$ from m' by increasing m' on one arbitrary node by $\nu\gamma - m'(V)$. Then m satisfies (3.9), (3.10), (3.11), and $\nu \mid m(V)$, thus by Theorem 3.14 there exists a ν-uniform hypergraph H with degree-vector m that covers p. The choice of m implies that H has γ hyperedges. □

3.4.4 k-edge-connectivity augmentation

As we have already seen, the k-edge-connectivity augmentation of an initial hypergraph $H_0 = (V, \mathcal{E}_0)$ corresponds to the covering of the set function $p(X) := (k - d_{H_0}(X))^+$ ($\emptyset \neq X \subset V$). This set function is symmetric and positively crossing supermodular, so Theorem 3.29 is applicable; furthermore, condition (3.23) concerning p-full partitions can be considerably simplified, as in Theorem 3.11. Recall that for a hypergraph $H = (V, \mathcal{E})$, $i_H(X) = |\{e \in \mathcal{E} : e \subseteq X\}|$, and $c(H)$ denotes the number of components of H.

Corollary 3.30. *Let $H_0 = (V_0, \mathcal{E}_0)$ be a hypergraph, and $\nu \geq 2$ an integer. There is a ν-uniform hypergraph H with γ hyperedges such that $H_0 + H$ is k-edge-connected if and only if the following hold:*

$$\nu\gamma \geq tk - \sum_{i=1}^{t} d_{H_0}(X_i) \quad \text{for every subpartition } \{X_1, \dots, X_t\}, \tag{3.24}$$

$$\gamma \geq k - d_{H_0}(X) \quad \text{for every } X \subseteq V, \tag{3.25}$$

$$(\nu - 1)\gamma \geq c(H_0 - \mathcal{E}'_0) - 1 \quad \text{for every } \mathcal{E}'_0 \subseteq \mathcal{E}_0 \text{ for which } |\mathcal{E}'_0| = k - 1. \tag{3.26}$$

Proof. Compared to Theorem 3.29 (applied on the p defined above), the only difference is that (3.23) is replaced by condition (3.26). Its necessity follows from the fact that the components of $H_0 - \mathcal{E}'_0$ form a p-full partition if $c(H_0 - \mathcal{E}'_0) > \nu$ (if $2 \leq c(H_0 - \mathcal{E}'_0) \leq \nu$, then the condition simply requires $\gamma \geq 1$).

We prove sufficiency by showing that if a p-full partition $\mathcal{F} = \{V_1, \dots, V_l\}$ violates (3.23) while conditions (3.24) and (3.25) are satisfied, then an appropriate $\mathcal{E}'_0$ violates (3.26). Let $\mathcal{E}'_0$ be the set of hyperedges in $\mathcal{E}_0$ that enter at least one member of $\mathcal{F}$, and let $H'_0 := (V, \mathcal{E}'_0)$. Then $d_{H'_0}(U) \leq k-1$ if U is the union of some (not all) members of $\mathcal{F}$, since $\mathcal{F}$ is p-full. The partition $\mathcal{F}$ violates (3.23), so $(\nu - 1)\gamma < l - 1 \leq c(H_0 - \mathcal{E}'_0) - 1$.

We claim that $|\mathcal{E}'_0| = k-1$, which then implies that (3.26) is violated. By (3.24), $lk - \sum_{i=1}^{l} d_{H'_0}(V_i) = lk - \sum_{i=1}^{l} d_{H_0}(V_i) \leq \nu\gamma \leq \frac{\nu}{\nu-1}(l-2) \leq 2l-4$, from which $\sum_{i=1}^{l} d_{H'_0}(V_i) \geq (k-2)l + 4$. This implies that $|\mathcal{E}'_0| \geq k-1$, and there are at least 4 members of $\mathcal{F}$ (say V_1, V_2, V_3, V_4), for which $d_{H'_0}(V_i) = k-1$. We can assume that $i_{H'_0}(V_1 \cup V_2) \leq i_{H'_0}(V_i \cup V_j)$ for every distinct $i, j \in \{1,2,3,4\}$.

If $i_{H'_0}(V_1 \cup V_2) > 0$, then $d_{H'_0}(V_1 \cup V_2) \geq d_{H'_0}(V_1) - i_{H'_0}(V_1 \cup V_2) + i_{H'_0}(V_2 \cup V_3) + i_{H'_0}(V_2 \cup V_4) \geq k$, contradicting the p-fullness of $\mathcal{F}$. So $i_{H'_0}(V_1 \cup V_2) = 0$, in which case there are $k-1$ hyperedges in $\mathcal{E}'_0$ that enter each of V_1, V_2, and $V_1 \cup V_2$. Suppose that $\mathcal{E}'_0$ contains a hyperedge besides these $k-1$, which enters a partition member V_i. Then $d_{H'_0}(V_1 \cup V_i) \geq k$, which would contradict the p-fullness of $\mathcal{F}$. This proves that $|\mathcal{E}'_0| = k-1$. □

3.4.5 Algorithmic aspects

It might be argued that Theorems 3.14 and 3.29 are not good characterizations, since it is not possible to check in polynomial time whether a given partition is p-full, hence it cannot be decided whether it is a deficient partition or not. Indeed, it is well known that it is NP-complete to decide for given graph $G = (V, E)$ whether there is a set $X \subseteq V$ with $d_G(X) \geq k$. Let $\mathcal{F}$ be the partition of V composed of singleton members, and let $p(X) := (k - d_G(X))^+$ for $\emptyset \neq X \subset V$ and $p(\emptyset) = p(V) := 0$. Now $\mathcal{F}$ is a p-full partition if and only if $d_G(X) < k$ for every $X \subseteq V$.

If a partition has at least one member V_i with $p(V_i) = 1$, then its deficiency can be checked using the characterization in Claim 3.17. But in general, deciding whether a partition is p-full or not is NP-complete. However, it is easy to see that if $p(V_i) \geq 2$ for every member of a deficient partition, then at least one of the partition members violates (3.9) (and the partition violates (3.21) in case of Theorem 3.29). This means that Theorems 3.14 and 3.29 give good co-NP characterizations.

The proof of Theorem 3.14 provides a polynomial algorithm for the degree specified problem if maximizing oracles are available for every set function of the form $p - d_H$, where H is an arbitrary hypergraph. In this case, a feasible splitting-off operation can be found in polynomial time using the procedure described in the proof. However, we have to introduce a kind of "multiple splitting-off" to ensure that the number of splitting-off steps is polynomial.

We choose a hyperedge e that can be feasibly split off by the method described in the proof of Theorem 3.14. If $m(v) \geq \nu$ for every $v \in e$ then the following *extended splitting-off operation* may be used:

($\star$) We determine the maximal μ for which e can be feasibly split off μ times and $\mu\chi_e(v) \leq m(v) - \nu + 1$ for every $v \in e$.

Claim 3.31. *An extended splitting-off can be executed in polynomial time.*

Proof. The degree restriction ensures that condition (3.17) is indifferent when calculating the number of feasible splitting-off operations, since the auxiliary graph is the same before each splitting-off. So it suffices to determine the maximal value μ for which the following hold:

$$\mu\chi_e(v) \leq m(v) - (\nu - 1) \quad \text{for every } v \in e, \tag{3.27}$$

$$\mu(|e \cap X| - 1) \leq m(X) - p(X) \quad \text{for every } X \text{ entered by } e, \tag{3.28}$$

$$\mu\nu \leq m(X) - p(X) \quad \text{if } e \subseteq X, \tag{3.29}$$

$$\mu \leq \frac{m(V)}{\nu} - p(X) \quad \text{if } e \subseteq X. \tag{3.30}$$

The maximal μ for which (3.29) and (3.30) hold can be determined by the maximizing oracle. As for (3.28), it suffices to check its validity on the family

$$\{X \subseteq V : \; p(Y) < p(X) \; \forall Y \subset X, \; e \text{ enters } X\}.$$

This family can be determined in polynomial time since it is laminar. □

The extended splitting-off operation is not used if $m(v) \leq \nu - 1$ for some $v \in e$; a single splitting-off is executed instead. But the number such splitting-off operations is polynomial (we can assume that $\nu \leq |V|$).

We have to prove that the number of extended splitting-off operations is also polynomial. First, observe that no set is deleted from $\mathcal{B}_1$ during an extended splitting-off. Indeed, $X \in \mathcal{B}_1$ is deleted during a single splitting-off of e if there is $Y \subset X$ for which $p(Y) < p(X)$ and $p^e(Y) = p^e(X)$; but then $|e \cap (X - Y)| > 0$, so $m^e(X - Y) \geq \nu - 1$, implying that $m^e(Y) - p^e(Y) \leq m^e(X) - (\nu - 1) - p^e(X) < 0$, which is impossible. Since $\mathcal{B}_1$ is always laminar, this implies that in a sequence of consecutive extended splitting-off operations, $\mathcal{B}_1$ can change only polynomially many times. It is easy to see that $\mathcal{B}_3$ also can change only polynomially many times. So it is enough to prove that if $\mathcal{B}_1$ and $\mathcal{B}_3$ do not change, then only polynomially many consecutive extended splittings are possible. But this follows from the fact that $\nu - 1$ nodes are the same in the hyperedges that are split off.

To solve the minimum cardinality problem in polynomial time we have to find a minimal degree specification m that satisfies (3.9). This can be done in polynomial time using the maximizing oracle. In the case of k-edge-connectivity augmentation, the required oracles can be realized using network flow algorithms.

Chapter 4

Connectivity augmentation of directed hypergraphs

In Section 1.3 it was shown that edge-connectivity of directed hypergraphs can be described in similar terms as edge-connectivity of digraphs. Consequently, the tools available for solving to edge-connectivity problems are also similar; however, there are some additional difficulties that must be overcome.

In this chapter we discuss edge-connectivity augmentation problems for directed hypergraphs. As in the undirected case, there are different ways to generalize the basic degree-specified and minimum cardinality digraphs problems. One possibility is minimum cardinality augmentation with restriction on the sizes of the new hyperarcs. Similarly to the undirected case, the problem with no restrictions is not interesting, since one could always choose hyperarcs containing every node of V. Unlike the undirected case, however, the augmentation with digraph edges is not more difficult than the digraph augmentation problem. The other possible objective is to minimize the total size of the added hyperarcs.

The main tool used in this chapter is a slight generalization of the method developed in [7] for splitting off in directed hypergraphs. This will be described in Section 4.2, after a brief introduction on augmentation with digraph edges. The chapter contains joint results with Márton Makai that appeared in [52].

4.1 Adding digraph edges

For a given directed hypergraph $D_0 = (V, \mathcal{A}_0)$, let $p(X) := (k - \varrho_{D_0}(X))^+$ for every $\emptyset \neq X \subset V$, and $p(\emptyset) = p(V) := 0$. Then the k-edge-connectivity augmentation of D with a minimum number of digraph edges corresponds to the problem of covering p with a minimum number of digraph edges. The set function p is positively crossing supermodular, so the following theorem of Frank [24] can be applied to solve the degree-constrained problem:

Theorem 4.1 (Frank [24]). *Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function; let $m_i : V \to \mathbb{Z}_+$ be an indegree-specification and $m_o : V \to \mathbb{Z}^+$ an outdegree-specification such that $m_i(V) = m_o(V)$. There exists a digraph $D = (V, A)$ such that $\varrho_D(v) = m_i(v)$ and $\delta_D(v) = m_o(v)$ $\forall v \in V$, and $\varrho_D(X) \geq p(X)$ $\forall X \subseteq V$ if and only if*

$$m_i(X) \geq p(X) \text{ for every } X \subseteq V \ ,$$

and

$$m_o(V - X) \geq p(X) \text{ for every } X \subseteq V \ . \quad \square$$

It should be noted that for the p defined above this result can be obtained from the directed splitting-off theorem of Mader [59], which is generalized by the above result of Frank.

The degree-specified result easily implies the following theorem on minimum cardinality augmentation:

Theorem 4.2 (Frank [24]). *Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function, and γ a non-negative integer. There exists a digraph $D = (V, A)$ with γ edges that covers p if and only if*

$$\gamma \geq \sum_{Z \in \mathcal{F}} p(X) \quad \text{for every partition } \mathcal{F} \text{ of } V,$$

and

$$\gamma \geq \sum_{Z \in \mathcal{F}} p(V - X) \quad \text{for every partition } \mathcal{F} \text{ of } V. \quad \square$$

The minimum cost k-edge-connectivity augmentation problem for digraphs is NP-complete even for $k = 1$, since it includes the problem of determining whether a digraph has a Hamiltonian cycle.

For undirected graphs, the local edge-connectivity augmentation problem was solvable [29] using Theorem 3.3 of Mader on splitting-off preserving local edge-connectivity. However, no such result can be expected for digraphs: local edge-connectivity augmentation of directed graphs is NP-complete, even in the special case when the requirement is 1 between the nodes of a specified subset T and 0 otherwise (see [29]).

4.2 Splitting-off in directed hypergraphs

In [7], Berg, Jackson and Jordán proved an interesting splitting-off theorem for directed hypergraphs, which led to a solution for the problem of directed hypergraph k-edge-connectivity augmentation by uniform hyperarcs.

In this section, we show that their edge-splitting result can be formulated in a more general form (using essentially the same proof). The result gives a method for solving a broader class of undirected and directed augmentation problems. The notion of splitting-off is used here in an abstract sense, similarly to Theorem 4.5 or Theorem 3.14.

4.2.1 Splitting off a single hyperarc

Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function. Let furthermore $m_i : V \to \mathbb{Z}_+$ be an indegree-specification and $m_o : V \to \mathbb{Z}_+$ an outdegree-specification for which $m_i(V) \leq m_o(V)$. Suppose that $m_i(X) \geq p(X)$ and $m_o(V - X) \geq p(X)$ for every $X \subseteq V$. We define the *splitting-off operation* analogously to the undirected definition in the proof of Theorem 3.14.

A hyperarc a can be *split off* from (p, m_i, m_o) if $m_i(h(a)) > 0$ and $\chi_{t(a)}(v) \leq m_o(v)$ for every $v \in V$. For such a hyperarc let

$$m_i^a(v) := \begin{cases} m_i(v) - 1 & \text{if } v = h(a), \\ m_i(v) & \text{otherwise,} \end{cases}$$

$$m_o^a(v) := m_o(v) - \chi_{t(a)}(v),$$

$$p^a(X) := \begin{cases} (p(X) - 1)^+ & \text{if } a \text{ enters } X, \\ p(X) & \text{otherwise.} \end{cases}$$

The splitting-off operation is *feasible* if $m_i^a(X) \geq p^a(X)$ and $m_o^a(V - X) \geq p^a(X)$ for every $X \subseteq V$. The operation is called a *feasible* $(r, 1)$-*splitting* if $|a| = r + 1$. It is easy to check using (1.4) that p^a is positively crossing supermodular.

To motivate the name "splitting-off", we may observe that if we add a new node z to V, and draw $m_i(v)$ digraph edges from z to each node $v \in V$, and $m_o(v)$ edges from each node $v \in V$ to z, then the above defined splitting-off operation corresponds to splitting-off in this digraph D. If in addition p is defined as $p(X) = (k - \varrho_{D_0}(X))^+$ for some digraph D_0, a feasible splitting-off corresponds to a splitting-off in $D_0 + D$ that preserves k-edge-connectivity in V (see Figure 4.1).

The following theorem describes conditions when a feasible splitting-off is available.

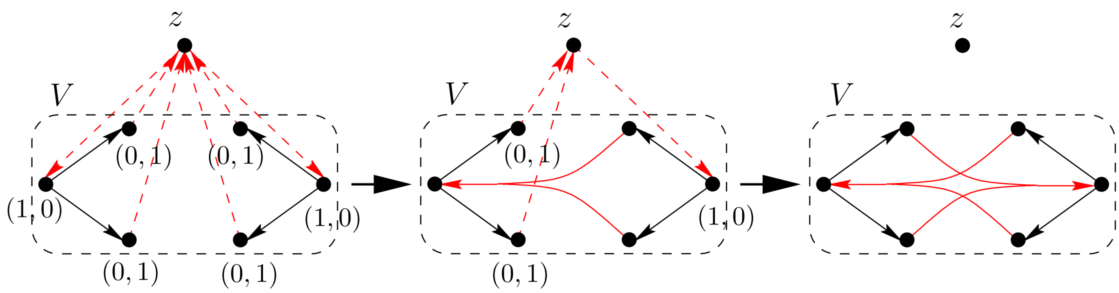

Figure 4.1: Two directed splitting-off steps preserving 1-edge-connectivity. The values at the nodes represent (m_i, m_o).

Theorem 4.3. *Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function, $m_i : V \to \mathbb{Z}_+$ and $m_o : V \to \mathbb{Z}_+$ degree specifications such that $m_i(V) \le m_o(V) \le rm_i(V)$ for some integer r, and*

$$m_i(X) \ge p(X) \text{ for every } X \subseteq V, \tag{4.1}$$

$$m_o(V - X) \ge p(X) \text{ for every } X \subseteq V. \tag{4.2}$$

Let $u \in V$ be such that $m_i(u) > 0$. Then there is a hyperarc a with $h(a) = u$ and $|a| \le r+1$ that can be feasibly split off.

Proof. We can assume that $m_i(V) \ge 2$. A set X is called *in-critical* if $u \in X$ and $p(X) = m_i(X)$. The maximal in-critical sets are pairwise co-disjoint, since they intersect, and by the crossing supermodularity of p, the union of two crossing in-critical sets is in-critical. The complement of a maximal in-critical set is called a *petal*. Let $\mathcal{F}$ denote the family of maximal in-critical sets, and let $\alpha := |\mathcal{F}|$; $\mathcal{F}$ is called an *α-flower*.

Claim 4.4. *$\alpha \le r$.*

Proof. Otherwise we would have

$$\sum_{X \in \mathcal{F}} p(X) = \sum_{X \in \mathcal{F}} m_i(X) > rm_i(V) \ge m_o(V) \ge \sum_{X \in \mathcal{F}} m_o(V - X),$$

which contradicts (4.2). □

First, suppose that $\alpha = 1$ ($a = \{u\}$ is obviously good for $\alpha = 0$), and let P be the single petal; $m_o(P) \ge p(V - P) = m_i(V - P) > 0$. A set X is called *out-critical* if $u \notin X$ and $m_o(V - X) = p(X) > 0$; if there are no such sets, then for any $v \in P$ with $m_o(v) > 0$ the digraph edge $a = vu$ can be split off. By the crossing supermodularity of p, the non-empty intersection of two out-critical sets is also out-critical. Since $u \notin X$ and

$m_i(X) \geq m_o(V-X)$ for any out-critical set, there are no two disjoint out-critical sets because of $m_i(V) \leq m_o(V)$, so there is a unique minimal out-critical Y. One of $P-Y$ and $Y-P$ is empty, otherwise $m_o(P-Y)+m_i(Y-P) < m_o(V-Y)+m_i(V-P) = p(Y)+p(V-P) \leq p(Y-P)+p(V-(P-Y)) \leq m_i(Y-P)+m_o(P-Y)$ would be a contradiction. Also, $m_o(Y) = m_o(V) - m_o(V-Y) \geq m_i(V) - m_o(V-Y) \geq m_i(V) - m_i(Y) > 0$, hence $m_o(Y \cap P) > 0$ follows from the emptiness of $P-Y$ or $Y-P$. Let $v \in Y \cap P$ be a node with $m_o(v) > 0$. Then the directed edge $a = vu$ can be feasibly split off.

If $\alpha \geq 2$, we define a by selecting as tail nodes one arbitrary node v with $m_o(v) > 0$ from each petal. We prove that a can be feasibly split off. By the construction of a, (4.1) holds after the splitting.

Suppose that there is a set X which violates (4.2) after the splitting, i.e. $m_o^a(V-X) < p^a(X)$. This means that if a enters X, then $p(X)-1 > m_o(V-X) - |a \cap (V-X)|$ while if a does not enter X, then $u \notin X$ and $p(X) > m_o(V-X) - |t(a) \cap (V-X)|$. In both cases $p(X) > m_o(V-X) - |a \cap (V-X)| + 1$.

There is a petal P such that $P - X \neq \emptyset$ and $X - P \neq \emptyset$ (this is trivial if X is subset of a petal; if it is not, then any petal P is good for which $P \cap a \notin X$, and such a petal exists otherwise $m_o(V-X) = m_o^a(V-X) < p^a(X) \leq p(X)$ contradicting (4.2)). The crossing supermodularity of p implies that

$$\begin{aligned} m_i(V-P) + m_o(V-X) - |a \cap (V-X)| + 1 < p(V-P) + p(X) \leq \\ \leq p(X-P) + p(V-(P-X)), \end{aligned}$$

so

$$\begin{aligned} m_i(X-P) + m_o(P-X) - |a \cap (P-X)| + 1 < p(X-P) + p(V-(P-X)), \\ m_o(P-X) - |a \cap (P-X)| + 1 < p(V-(P-X)), \end{aligned}$$

which would imply that $V-(P-X)$ violates (4.2), since $|a \cap (P-X)| \leq 1$. □

The theorem implies that there is a hyperedge a that can be feasibly split off such that in addition $m_i^a(V) \leq m_o^a(V)$ holds. This guarantees that subsequent splitting-off operations are possible, eventually leading to a complete splitting-off. In the following we describe the consequences.

4.2.2 Complete splitting-off

The next theorem states that a complete splitting-off of uniform hyperedges is always possible if the obvious necessary conditions hold. A special case (when for a given directed hypergraph D and positive integer k, $p(X) = (k - \varrho_D(X))^+$ for every $\emptyset \neq X \subset V$) was proved in [7]. The more general result presented here has the advantage that it can be applied to extensions of k-edge-connectivity as well (while the proof is not more difficult).

Theorem 4.5. *Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function, $m_i : V \to \mathbb{Z}_+$ and $m_o : V \to \mathbb{Z}_+$ degree specifications such that $m_o(V) = rm_i(V)$ for some positive integer r, and*

$$m_i(X) \geq p(X) \text{ for every } X \subseteq V, \tag{4.3}$$

$$m_o(V - X) \geq p(X) \text{ for every } X \subseteq V. \tag{4.4}$$

Then there is a directed $(r,1)$-hypergraph D such that $\delta_D(v) = m_o(v)$ and $\varrho_D(v) = m_i(v)$ for every $v \in V$, and

$$\varrho_D(X) \geq p(X) \text{ for every } X \subseteq V.$$

Proof. According to Theorem 4.3 we can obtain a directed hypergraph D^* by successive feasible splitting-off operations such that $\delta_{D^*}(v) = m_o(v)$, $\varrho_{D^*}(v) = m_i(v)$ for every $v \in V$, and $\varrho_{D^*}(X) \geq p(X)$ for every $X \subseteq V$. Since $m_o(V) = rm_i(V)$, $|a| \geq r + 1$ holds for at least one hyperarc a of D^*. So there is a feasible $(r_1, 1)$-splitting with head $h(a)$ for some $r_1 \geq r$; moreover, by Theorem 4.3 there is also a feasible $(r_2, 1)$-splitting with head $h(a)$ for some $r_2 \leq r$.

Lemma 4.6. *If for some $r_1 > r > r_2$ there is a feasible $(r_1, 1)$-spitting and a feasible $(r_2, 1)$-splitting with head u , then there is a feasible $(r, 1)$-splitting with head u.*

Proof. Let a be the hyperarc obtained by the $(r_1, 1)$-splitting. By induction, it suffices to show that for some $v \in t(a)$, the hyperarc a' defined by $a' = a - v$, $h(a') = u$ gives a feasible splitting. If $\chi_a(v) \geq 2$ for some $v \in V$ then such a v is suitable, so we may assume that $\chi_a \leq 1$. If $r_1 > 2$, then suppose indirectly that for every $v \in t(a)$ there is an in-critical set X_v such that $a - v \subseteq X_v$ and $v \notin X_v$. We can assume that these are maximal in-critical sets. Thus the sets $\{X_v : v \in t(a)\}$ form a flower with r_1 petals centered on u; but this contradicts the fact that there is a feasible $(r_2, 1)$-splitting with head u.

If $r_1 = 2$, then $r_2 = 0$, so there are no in-critical sets. As we have seen in the proof of Theorem 4.3, there is a unique minimal out-critical set Y with $u \notin Y$. Then $a - Y = u$, otherwise the $(2, 1)$-splitting would not be feasible; thus both $(1, 1)$-splittings are feasible. ☐

We prove Theorem 4.5 by induction on $m_i(V)$. According to Lemma 4.6, there is a node $u \in V$ with $m_i(u) > 0$ for which there exists a feasible $(r, 1)$-splitting at u; let a be the resulting $(r,1)$-hyperarc. By induction, there is a directed $(r,1)$-hypergraph D' that satisfies the conditions given by m_i^a, m_o^a and p^a. The directed hypergraph obtained by adding a to D' satisfies the conditions of Theorem 4.5. □

In fact, a bit more general theorem can also be proved by the same method (this was also remarked by Berg, Jackson, and Jordán for the k-edge-connectivity case):

Theorem 4.7. *Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function, $m_i : V \to \mathbb{Z}_+$ and $m_o : V \to \mathbb{Z}_+$ degree specifications such that $m_i(V) \leq m_o(V)$ and $(r-1)m_i(V) < m_o(V) \leq rm_i(V)$ for some positive integer r, and*

$$m_i(X) \geq p(X) \text{ for every } X \subseteq V, \tag{4.5}$$

$$m_o(V - X) \geq p(X) \text{ for every } X \subseteq V. \tag{4.6}$$

Then there is a directed hypergraph D consisting of $(r,1)$-hyperedges and $(r-1,1)$-hyperedges such that $\delta_D(v) = m_o(v)$ and $\varrho_D(v) = m_i(v)$ for every $v \in V$, and

$$\varrho_D(X) \geq p(X) \text{ for every } X \subseteq V.$$

Proof. By Theorem 4.3 there is a complete splitting-off, so there must be an $(r_1,1)$-hyperarc a that can be feasibly split off for some $r_1 \geq r$, since $(r-1)m_i(V) < m_o(V)$. By Theorem 4.3 there is also an $(r_2,1)$-hyperarc for some $r_2 \leq r$ with head $h(a)$ that can be feasibly split off, since $m_o(V) \leq rm_i(V)$. So Lemma 4.6 implies that there is an $(r,1)$-hyperarc that can be feasibly split off. We can continue to split off $(r,1)$-hyperarcs as long as $(r-1)m_i'(V) < m_o'(V) \leq rm_i'(V)$ holds for the modified degree-specifications. It is easy to see that the first time this does not hold is when $(r-1)m_i'(V) = m_o'(V)$. But then by Theorem 4.5 we can finish the complete splitting-off by splitting off $(r-1,1)$-hyperarcs. □

4.3 Covering supermodular functions by directed hypergraphs

Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function. When covering p with a directed hypergraph, we can have various objectives just as in the case of undirected hypergraphs: we can minimize the total size of the added hyperarcs, or we can cover p with an $(r,1)$-uniform directed hypergraph of minimum cardinality.

4.3.1 Adding hyperarcs of minimum total size

Theorem 4.8. *Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function. The minimum total size σ of a directed hypergraph that covers p is $\sigma = \sigma_1 + \sigma_2$, where*

$$\sigma_1 = \max\left\{\max\left\{\sum_{Z\in\mathcal{F}} p(X), \left\lceil \frac{1}{|\mathcal{F}|-1}\left(\sum_{Z\in\mathcal{F}} p(V-X)\right)\right\rceil\right\} : \ \mathcal{F} \text{ a partition of } V\right\},$$

$$\sigma_2 = \max\left\{\max\left\{\sum_{Z\in\mathcal{F}} p(X), \sum_{Z\in\mathcal{F}} p(V-X)\right\} : \ \mathcal{F} \text{ a partition of } V\right\}.$$

Proof. To see the necessity of the conditions, let $D = (V, \mathcal{A})$ be a directed hypergraph of minimum total size that covers p, and let $m_i(v) := \varrho_D(v)$ and $m_o(v) := \delta_D(v)$ for every $v \in V$. Then $m_i(X) \geq p(X)$ and $m_o(V - X) \geq p(X)$ for every $X \subseteq V$. Let $\mathcal{F}$ be a partition of V. Then $m_i(V) \geq \sum_{Z\in\mathcal{F}} p(X)$, and $(|\mathcal{F}| - 1)m_i(V) \geq \sum_{Z\in\mathcal{F}} p(V - X)$. Similarly, $m_o(V) \geq \sum_{Z\in\mathcal{F}} p(V - X)$, and $(|\mathcal{F}| - 1)m_o(V) \geq \sum_{Z\in\mathcal{F}} p(X)$; in addition, we know that $m_o(V) \geq m_i(V)$. We obtained that $m_i(V) \geq \sigma_1$ and $m_o(V) \geq \sigma_2$, thus the total size of D is at least σ.

Sufficiency is proved by finding degree specifications m_i and m_o for which we can apply the complete splitting-off described in Theorem 4.7. So we have to find m_i satisfying (4.5) and m_o satisfying (4.6) such that $m_i(V) \leq m_o(V)$.

Let p_1 be the set function defined by $p_1(X) := p(X)$ if $X \subset V$ and $p_1(V) := \sigma_1$. If the conditions of the theorem hold, p_1 is positively crossing supermodular, and by Theorem 2.13 $B(p_1)$ is a non-empty base polyhedron. Let m_i be an integral vector in $B(p_1)$; then m_i satisfies (4.5).

Let p_2 be defined by $p_2(X) = p(V - X)$ if $\emptyset \neq X \subset V$, and $p_2(V) := \sigma_2$. If the conditions of the theorem hold, then p_2 is positively crossing supermodular, and $B(p_2)$ is a non-empty base polyhedron by Theorem 2.13. Let m_o be an integral vector in $B(p_2)$; then m_o satisfies (4.6) and $m_i(V) = \sigma_1 \leq \sigma_2 = m_o(V)$.

We can apply Theorem 4.7 on these m_i and m_o values, and a suitably chosen r. □

4.3.2 Covering by uniform directed hypergraphs

We can use the degree-specified result in Theorem 4.5 to give a formula on the minimal number of $(r, 1)$-hyperedges that can cover a given positively crossing supermodular set function.

Theorem 4.9. *Let $p : 2^V \to \mathbb{Z}_+$ be a positively crossing supermodular set function, and r a positive integer. There exists a directed $(r, 1)$-hypergraph with γ hyperarcs that covers p*

if and only if

$$\gamma \geq \sum_{X \in \mathcal{F}} p(X), \tag{4.7}$$

$$r\gamma \geq \sum_{X \in \mathcal{F}} p(V - X), \tag{4.8}$$

$$(|\mathcal{F}| - 1)\,\gamma \geq \sum_{X \in \mathcal{F}} p(V - X) \tag{4.9}$$

hold for every partition $\mathcal{F}$ of V.

Proof. The necessity of the conditions can be seen easily. To prove sufficiency, one can construct degree specifications m_i and m_o that satisfy the conditions of Theorem 4.5. Let us define the set function p' by $p'(X) := p(X)$ for $X \subset V$ and $p'(V) := \gamma$. Then p' is positively crossing supermodular. Since (4.7) and (4.9) hold for p', by applying Theorem 2.13 to p' we get a nonnegative integer vector m_i such that $m_i(X) \geq p(X)$ for all $X \subseteq V$, and $m_i(V) = \gamma$.

To construct m_o let m'_o be a nonnegative vector satisfying $m'_o(V - X) \geq p(X)$ for all $X \subseteq V$, which is minimal in the sense that for every $v \in V$ with $m'_o(v) > 0$ there exists a set X for which $v \notin X$ and $m'_o(V - X) = p(X)$ (such a set is called *tight*). Let $\mathcal{B} = \{X_1, X_2, \ldots, X_l\}$ be a family of minimum cardinality which for every node v with $m'_o(v) > 0$ contains a tight set not containing v. Then $\mathcal{B}$ is cross-free, since we could replace a crossing pair X_i, X_j by their intersection according to (2.2). If the family is composed of co-disjoint sets, then

$$m'_o(V) = \sum_{i=1}^{l} m'_o(V - X_i) = \sum_{i=1}^{l} p(X_i) \leq r\gamma$$

by (4.8). If there are two disjoint sets X_i and X_j in $\mathcal{B}$, then

$$m'_o(V) \leq m'_o(V - X_i) + m'_o(V - X_j) = p(X_i) + p(X_j) \leq \gamma.$$

Thus we can obtain an out-degree specification m_o by increasing m'_o on an arbitrary node to obtain $m_o(V) = r\gamma$. By Theorem 4.5 there is a directed $(r, 1)$-hypergraph with degrees m_i and m_o that covers p. □

The following example demonstrates that condition (4.9) cannot be left out. Let $V = \{v_1, v_2, v_3\}$, $p(\{v_1, v_2\}) = p(\{v_1, v_3\}) = p(\{v_2, v_3\}) := 2$ and $p(X) := 0$ for the other sets, $r := 3$, $\gamma := 2$. Conditions (4.7) and (4.8) are satisfied, but (4.9) is not, and there is no directed $(3, 1)$-hypergraph of 2 hyperarcs covering p. When $r = 1$, i.e. we add digraph edges, (4.9) follows from (4.8).

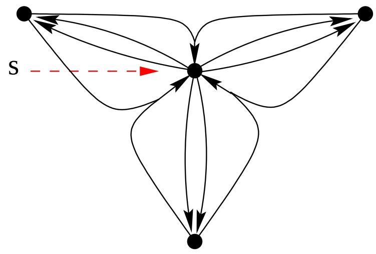

Figure 4.2: A directed hypergraph that is $(2,2)$-edge-connected from root s but cannot be decomposed into two directed sub-hypergraphs connected to s.

4.4 (k,l)-edge-connectivity of directed hypergraphs

In this section the definition of (k,l)-edge-connectivity, which was already introduced for graphs in Chapter 1, is extended to directed hypergraphs as a common generalization of k-edge-connectivity and k-rooted-connectivity. This connectivity property can be characterized using crossing supermodular set functions, so the tools developed in this chapter are applicable.

4.4.1 Properties

Let $D=(V,\mathcal{A})$ be a directed hypergraph, and $s \in V$ a specified root node. D is called (k,l)*-edge-connected* from root s for non-negative integers k and l if there there are k edge-disjoint paths from s to any other node, and there are l edge-disjoint paths to s from any other node. We say that D is (k,l)-edge-connected if there is a node $s \in V$ such that D is (k,l)-edge-connected from root s.

By Proposition 1.12, a directed hypergraph D is (k,l)-edge-connected from root s if and only if $\varrho_D(X) \geq k$ if $\emptyset \neq X \subset V-s$, and $\varrho_D(X) \geq l$ if $s \in X \subset V$. It is easy to see that D is (k,k)-edge-connected if and only if it is k-edge-connected, while D is $(k,0)$-edge-connected form root s if and only if it is k-rooted-connected from root s.

We have seen in Chapter 1 that by Proposition 1.14, a directed hypergraph that is k-rooted-connected from s can be decomposed into k edge-disjoint spanning directed sub-hypergraphs that are connected from s. As a consequence, the same is true if a directed hypergraph D is (k,l)-edge-connected from s. If D is a digraph, then it is also true that D can be decomposed into l edge-disjoint spanning sub-digraphs that are connected from $V-s$ to s. However, the latter claim is not true for directed hypergraphs, as the example on Figure 4.2 shows.

The proof of Theorem 1.14 in Chapter 1 implied that if a directed hypergraph $D=(V,\mathcal{A})$

is k-rooted-connected, then it can be shrinked (meaning that we delete all but one tail-nodes of every hyperarc) to a k-rooted connected digraph. For (k,l)-edge-connectivity no similar statement is true; in fact, Bang-Jensen and Thomassé [5] observed the following:

Proposition 4.10 ([5]). *It is NP-complete to decide whether a strongly connected directed hypergraph can be shrinked to a strongly connected digraph.*

Proof. Let $D=(V,A)$ be a strongly connected digraph. We define a directed hypergraph $D'=(V,\mathcal{A})$ the following way: $\mathcal{A}$ contains one hyperarc a_v for every $v\in V$, namely $h(a_v)=v$, and $t(a_v)=\{u\in V:\ uv\in A\}$. Then D' is strongly connected, and it can be shrinked to a strongly connected digraph if and only if D contains a Hamiltonian cycle. □

4.4.2 (k,l)-edge-connectivity augmentation

Let $D_0=(V,\mathcal{A}_0)$ be a directed hypergraph, and let k,l be non-negative integers. We are interested in the (k,l)-edge-connectivity augmentation of D_0, either with the objective of minimizing the total size of the new hyperarcs, or with the objective of adding the minimum number of $(r,1)$-hyperarcs . To solve these problems using the results in Section 4.3, we define the following set function:

$$p(X):=\begin{cases}(k-\varrho_{D_0}(X))^+ & \text{if } \emptyset\neq X\subseteq V-s,\\ (l-\varrho_{D_0}(X))^+ & \text{if } s\in X\subset V,\\ 0 & \text{if } X=\emptyset \text{ or } X=V.\end{cases}\tag{4.10}$$

The set function p is positively crossing supermodular, so we can apply Theorems 4.8 and 4.9. If $k\geq l$, then the conditions of the theorems can be simplified:

Claim 4.11. *Let $k\geq l$, and let p be the set function defined in (4.10). If*

$$\gamma\geq\sum_{Z\in\mathcal{F}}p(Z)$$

for every partition $\mathcal{F}$ of V, then

$$(|\mathcal{F}|-1)\,\gamma\geq\sum_{Z\in\mathcal{F}}p(V-Z)$$

for every partition $\mathcal{F}$.

Proof. Suppose that $(t-1)\gamma<\sum_{Z\in\mathcal{F}}p(V-Z)$ for a partition $\mathcal{F}=\{X_1,\dots,X_t\}$ of V with $t\geq 2$. If $p(V-X_t)=0$, then $\sum_{Z\in\mathcal{F}}p(V-Z)=\sum_{i=1}^{t-1}p(V-X_i)\leq\sum_{i=1}^{t-1}(p(X_i)+p(V-X_i))\leq(t-1)\gamma$, by considering the two-member partitions $\{X_i,V-X_i\}$. So we can assume that $p(V-X_i)>0$ for every i, thus $\sum_{Z\in\mathcal{F}}p(V-Z)=k+(t-1)l-\sum_{i=1}^{t}\delta_{D_0}(X_i)\leq k+(t-1)l-\sum_{i=1}^{t}\rho_{D_0}(X_i)\leq(t-1)k+l-\sum_{i=1}^{t}\rho_{D_0}(X_i)\leq\sum_{i=1}^{t}p(X_i)\leq\gamma$. □

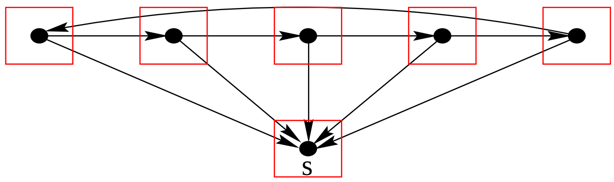

Figure 4.3: Counterexample for Claim 4.11 for $k = 1$, $l = 3$. Inequality (4.9) on the indicated copartition gives $\gamma \geq 2$, while $\sum_{Z\in\mathcal{F}} p(Z) \leq 1$ for every partition $\mathcal{F}$.

The counterexample on Figure 4.3 shows that the claim is not necessarily true if $k < l$.

Using Claim 4.11, Theorem 4.8 implies the following on minimum total size augmentation:

Theorem 4.12. *Let $D_0 = (V, \mathcal{A}_0)$ be a directed hypergraph, let $k \geq l$ be non-negative integers, and let p be the set function defined in (4.10). The minimum total size σ of a directed hypergraph D for which $D_0 + D$ is (k,l)-edge-connected from root s is $\sigma = \sigma_1 + \sigma_2$, where*

$$\sigma_1 = \max\left\{\sum_{Z\in\mathcal{F}} p(X) : \ \mathcal{F} \text{ is a partition of } V\right\},$$
$$\sigma_2 = \max\left\{\max\left\{\sum_{Z\in\mathcal{F}} p(X), \sum_{Z\in\mathcal{F}} p(V-X)\right\} : \ \mathcal{F} \text{ is a partition of } V\right\}.$$

□

Theorem 4.13. *Let $D_0 = (V, \mathcal{A}_0)$ be a directed hypergraph, let $k < l$ be non-negative integers, and let p be the set function defined in (4.10). The minimum total size σ of a directed hypergraph D for which $D_0 + D$ is (k,l)-edge-connected from root s is $\sigma = \sigma_1 + \sigma_2$, where*

$$\sigma_1 = \max\left\{\max\left\{\sum_{Z\in\mathcal{F}} p(X), \left\lceil \frac{1}{|\mathcal{F}|-1}\Big(\sum_{Z\in\mathcal{F}} p(V-X)\Big)\right\rceil\right\} : \ \mathcal{F} \text{ is a partition of } V\right\},$$
$$\sigma_2 = \max\left\{\max\left\{\sum_{Z\in\mathcal{F}} p(X), \sum_{Z\in\mathcal{F}} p(V-X)\right\} : \ \mathcal{F} \text{ is a partition of } V\right\}.$$

□

For minimum cardinality augmentation using $(r,1)$-hyperedges, the following results are obtained from Theorem 4.9 and Claim 4.11:

Theorem 4.14. *Let $D_0 = (V, \mathcal{A}_0)$ be a directed hypergraph, and let $k \geq l$ and $r \geq 1$ be non-negative integers. D_0 can be made (k,l)-edge-connected by the addition of γ new $(r,1)$-hyperarcs if and only if*

$$\begin{aligned} \gamma &\geq \sum_{X \in \mathcal{F}} p(X), \\ r\gamma &\geq \sum_{X \in \mathcal{F}} p(V - X) \end{aligned}$$

hold for every partition $\mathcal{F}$ of V, where p is the set-function defined in (4.10).

Theorem 4.15. *Let $D_0 = (V, \mathcal{A}_0)$ be a directed hypergraph, and let $k < l$ and $r \geq 1$ be non-negative integers. D_0 can be made (k,l)-edge-connected by the addition of γ new $(r,1)$-hyperarcs if and only if*

$$\begin{aligned} \gamma &\geq \sum_{X \in \mathcal{F}} p(X), \\ r\gamma &\geq \sum_{X \in \mathcal{F}} p(V - X), \\ (|\mathcal{F}| - 1)\,\gamma &\geq \sum_{X \in \mathcal{F}} p(V - X) \end{aligned}$$

hold for every partition $\mathcal{F}$ of V, where p is the set-function defined in (4.10).

Chapter 5

Partition-connectivity of hypergraphs

5.1 Introduction

In Chapter 1, the notion of k-partition-connectivity was introduced for graphs. In this chapter, we discuss an extension of this notion to hypergraphs. While 1-partition-connectivity for graphs is simply equivalent to connectivity, the extension for hypergraphs defines a different type of connectivity property. In this section we define this extension and describe some of its basic characteristics. The results of the chapter are joint results with András Frank and Matthias Kriesell [37].

5.1.1 A different approach to connectivity

Recall that an equivalent characterization of connectivity of graphs is that in order to dismantle the graph into $t+1$ components, one has to delete at least t edges. The k-partition-connectivity of graphs can also be defined from this point of view: in order to dismantle the graph into $t+1$ components, one has to delete at least kt edges. We start with the introduction of this type of characterization for hypergraphs..

A hypergraph $H = (V, \mathcal{E})$ is called *partition-connected* if one has to delete at least t hyperedges to dismantle it into $t+1$ components for every t. An equivalent formulation is that $e_H(\mathcal{F}) \geq |\mathcal{F}| - 1$ holds for every partition $\mathcal{F}$ of V. Partition-connectivity clearly implies connectivity; contrary to the graph case, however, a connected hypergraph is not necessarily partition-connected (a partition-connected hypergraph must have at least $|V|-1$ hyperedges).

Theorem 1.5 of Tutte characterized graphs that can be decomposed into k connected spanning subgraphs. In light of the previous remarks, at least two different generalizations are possible for hypergraphs. We say that a sub-hypergraph $(V, \mathcal{E}')$ of a hypergraph

$H = (V, \mathcal{E})$ is *spanning* if $V = \cup(\mathcal{E}')$. One can consider the problem of decomposing a hypergraph into k connected spanning sub-hypergraphs, or the problem of decomposition into k partition-connected spanning sub-hypergraphs. As the following theorem shows, the former problem is NP-complete, while we will see that the latter is solvable and has relatively simple structural properties.

Theorem 5.1. *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 $k \geq 2$.*

Proof. First assume that $k = 2$. Recall that it is NP-complete to decide whether the nodes of a hypergraph can be coloured by 2 colours such that every hyperedge contains nodes of both colours. This implies by duality that colouring the hyperedges of a hypergraph $H_0 = (V_0, \mathcal{E}_0)$ by red and blue so that every node belongs to both a red and a blue hyperedge is also NP-complete. We show that this latter problem is polynomially solvable if there is a polynomial algorithm to decide decomposability of a hypergraph into two connected spanning sub-hypergraphs. Let s be a new node, let $V := V_0 + s$, $\mathcal{E} := \{e + s : e \in \mathcal{E}_0\}$, and $H = (V, \mathcal{E})$. Note that a sub-hypergraph $H' = (V, \mathcal{E}')$ of H is connected and spans V if and only if the corresponding sub-hypergraph $(V_0, \mathcal{E}'_0)$ of H_0 spans V_0. Therefore H can be decomposed into 2 connected spanning sub-hypergraphs if and only if H_0 can be decomposed into two sub-hypergraphs both spanning V_0.

The NP-completeness of the problem for $k \geq 3$ easily reduces to the case $k = 2$. Let $H = (V, \mathcal{E})$ be a hypergraph and let H^+ denote the hypergraph arising from H by adding $k - 2$ copies of V as new hyperedges. Then H^+ can be decomposed into k connected spanning sub-hypergraphs if and only if H can be decomposed into 2 connected spanning sub-hypergraphs. □

5.1.2 k-partition-connectivity

A hypergraph $H = (V, E)$ is called *k-partition-connected* if for every t, one has to delete at least kt hyperedges to dismantle it into $t+1$ components. Equivalently, $e_H(\mathcal{F}) \geq k(|\mathcal{F}|-1)$ for every partition $\mathcal{F}$ of V. This is clearly a generalization of k-partition-connectivity for graphs. In the graph case, the property was also definable in terms of decomposition into connected spanning subgraphs (see Theorem 1.5). As an analogous result, we will show that k-partition-connectivity of a hypergraph is equivalent to decomposability into k partition-connected spanning sub-hypergraphs.

Another equivalent characterization of k-partition-connectivity of graphs is the existence of k edge-disjoint spanning trees. The problem of finding k edge-disjoint spanning trees in a graph is a special case of finding k disjoint bases of a matroid and therefore Theorem

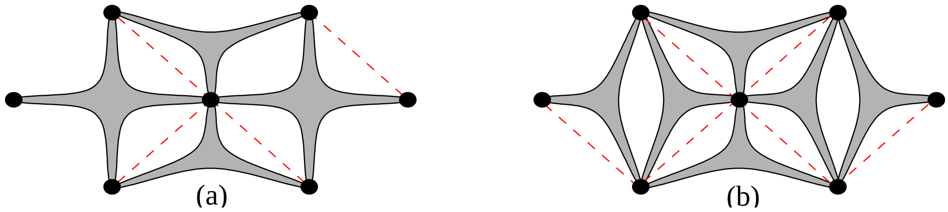

Figure 5.1: (a) a wooded hypergraph; (b) a tree-reducible hypergraph

1.5 may be easily derived from Edmonds' matroid partition theorem (Theorem 2.7). Our approach for hypergraphs also makes use of this result and is based on an observation of Lorea [54] that the notion of circuit matroids of graphs can be generalized to hypergraphs. We elaborate this in the next section, by introducing the structures that play a role similar to spanning trees in graphs.

5.2 Tree-reducible hypergraphs

We call a hypergraph $H = (V, \mathcal{E})$ *forest-reducible* or *wooded* if it is possible to select two nodes from each hyperedge of H so that the chosen pairs, as graph edges, form a forest (Figure 5.1). If this forest may be chosen to be a spanning tree, then H is called *tree-reducible*. Clearly a tree-reducible hypergraph is partition-connected. We will prove that every partition-connected hypergraph contains a tree-reducible sub-hypergraph, and the role of these sub-hypergraphs is similar to the role of spanning trees of graphs with respect to k-partition-connectivity. First we need some additional definitions.

5.2.1 Hyperforests and the hypergraphic matroid

Given a hypergraph $H = (V, \mathcal{E})$, we can define $b : 2^{\mathcal{E}} \to \mathbb{Z}_+$ as $b(\mathcal{X}) := |\cup(\mathcal{X})|$ for every $\mathcal{X} \subseteq \mathcal{E}$. It is easy to see that b is fully submodular on the ground set $\mathcal{E}$. We say that the *strong Hall condition* holds for the hypergraph H if $|\cup(\mathcal{X})| \geq |\mathcal{X}| + 1$ for every $\emptyset \neq \mathcal{X} \subseteq \mathcal{E}$.

It should be remarked that the notion comes from the theory of bipartite graphs: a bipartite graph $G = (U, V; E)$ satisfies the strong Hall condition for U if $|\Gamma(X)| \geq |X| + 1$ for every $\emptyset \neq X \subseteq U$, where $\Gamma(X) = \{v \in V : \ uv \in E \text{ for some } u \in X\}$. There is a standard way of associating a bipartite graph $G = (U, V; E)$ to a hypergraph $H = (V, \mathcal{E})$ by setting $U := \mathcal{E}$, and for $e \in U$ and $v \in V$, $ev \in E$ if and only if $v \in e$. Thus the strong Hall condition for the hypergraph corresponds to the strong Hall condition for the associated bipartite graph.

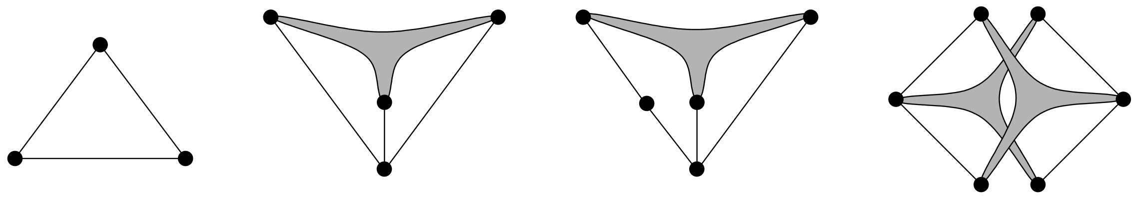

Figure 5.2: Some examples of hypercircuits

A bipartite graph $G = (U, V; E)$ is called *elementary* if G is connected, $|U| = |V|$, and $\Gamma(X) \geq |X| + 1$ holds for every $\emptyset \neq X \subset U$ (which is equivalent to requiring the inequality for nonempty proper subsets of V.) By a *hypercircuit* we mean a hypergraph the associated bipartite graph of which is elementary (see Figure 5.2). Note that in the special case when the hypergraph is a graph, this notion coincides with the usual notion of a graph circuit. By a *hyperforest* we mean a hypergraph $H = (V, \mathcal{E})$ where there is no subset $\mathcal{X} \subseteq \mathcal{E}$ for which $H' = (\cup(\mathcal{X}), \mathcal{X})$ is a hypercircuit. This is equivalent to saying that H satisfies the strong Hall condition, or that there are at most $|X| - 1$ hyperedges of H induced by X for every $\emptyset \neq X \subseteq V$. A hyperforest $H = (V, \mathcal{E})$ is a *spanning hypertree* if $\cup(\mathcal{E}) = V$ and $|\mathcal{E}| = |V| - 1$.

Erdős had conjectured and Lovász [56] proved that the node set of a hyperforest can always be coloured by two colours so that every hyperedge contains nodes of both colours. What Lovász actually proved (in a more general form) was the following result (which can also be derived from Theorem 2.8 of Edmonds on matroid intersection):

Theorem 5.2. *A hypergraph H is a hyperforest if and only if it is wooded.*

This result clearly implies Erdős' conjecture since the hypergraph can be reduced to a forest that is bipartite, and a two-colouring of its nodes gives the required two-colouring of the hypergraph. It should be mentioned that the proof of Lovász is constructive, so it gives an algorithm for deciding whether a hypergraph is wooded. (Note that the word *wooded* in Hungarian translates to *erdős*.)

Let us turn our attention to the relation between wooded sub-hypergraphs and matroid theory. In Chapter 1 we already mentioned the well-known fact that the subforests of an undirected graph $G = (V, E)$ form the family of independent sets of a matroid on the ground set E, called the circuit matroid of G. Lorea [54] extended this notion to hypergraphs.

Theorem 5.3 (Lorea). *Given a hypergraph $H = (V, \mathcal{E})$, the sub-hypergraphs of H which are hyperforests (or equivalently, the wooded sub-hypergraphs of H) form the family of independent sets of a matroid on the ground set $\mathcal{E}$.*

Proof. We define the set-function $b_H : 2^{\mathcal{E}} \to \mathbb{Z}_+$ by

$$b_H(\mathcal{X}) = |\cup(\mathcal{X})| - 1 \tag{5.1}$$

for a non-empty subset $\mathcal{X} \subseteq \mathcal{E}$, and $b_H(\emptyset) := 0$. Since $|\cup(\mathcal{X})|$ is a fully submodular function of the ground set $\mathcal{E}$, b_H is intersecting submodular, and it is obviously monotone increasing. Let us consider the matroid M_{b_H} defined in Theorem 2.16. In this a subset $\mathcal{X} \subseteq \mathcal{F}$ is independent if $|\mathcal{Z}| \leq b_H(\mathcal{Z})$ holds for every subset $\mathcal{Z} \subseteq \mathcal{X}$. This is equivalent to requiring that the sub-hypergraph $H' = (V, \mathcal{X})$ should meet the strong Hall condition, that is, by Theorem 5.2, H' should be wooded. □

A matroid arising this way is called the *circuit matroid* of the hypergraph H and it is denoted by M_H. We call a matroid which is the circuit matroid of a hypergraph a *hypergraphic matroid*. By choosing H to be the hypergraph consisting of a three-element ground set V and of four copies of V as hyperedges, we can observe that the uniform matroid $U_{4,2}$ (the smallest non-binary matroid) is hypergraphic. It should be noted that unlike graphic matroids, the class of hypergraphic matroids is not closed under contraction.

5.2.2 Rank function

For a hypergraph $H = (V, \mathcal{E})$, let b_H be the set function defined by (5.1) for $\emptyset \neq \mathcal{X} \subseteq \mathcal{E}$. We can apply Theorem 2.16 of Edmonds to determine the rank-function of the circuit matroid M_H.

Theorem 5.4. *The rank function r_H of the circuit matroid of a hypergraph H is given by the following formula:*

$$r_H(\mathcal{X}) = \min\{|V| - |\mathcal{F}| + e_{\mathcal{X}}(\mathcal{F}) : \ \mathcal{F} \text{ is a partition of } V\}. \tag{5.2}$$

Proof. It suffices to prove the formula for the special case $\mathcal{X} = \mathcal{E}$ since the value of $e_{\mathcal{X}}(\mathcal{F})$ does not change if the hyperedges in $\mathcal{E} - \mathcal{X}$ are deleted.

Let $H' = (V, \mathcal{E}')$ be a wooded sub-hypergraph of H. Then, for every partition $\mathcal{F}$ of V, there are at most $|V| - |\mathcal{F}|$ hyperedges in H' which are induced by a member of $\mathcal{F}$ according to the strong Hall condition. Therefore $|\mathcal{E}'|$ cannot be bigger than $|V| - |\mathcal{F}| + e_{\mathcal{E}'}(\mathcal{F})$, that is, $r_H(\mathcal{E}) \leq |V| - |\mathcal{F}| + e_{\mathcal{E}}(\mathcal{F})$. We have to prove the existence of a partition for which equality holds.

By Theorem 2.16,

$$r_H(\mathcal{E}) = \min\{\sum_{i=1}^{t} b_H(\mathcal{Z}_i) + |\mathcal{E} - (\cup_{i=1}^{t} \mathcal{Z}_i)| : \ \{\mathcal{Z}_1, \ldots, \mathcal{Z}_t\} \text{ is a subpartition of } \mathcal{E}\}. \tag{5.3}$$

A subpartition where the minimum is attained will be called a *minimizer* of (5.3). Let $\mathcal{N} := \{\mathcal{Z}_1, \ldots, \mathcal{Z}_t\}$ be a minimizer of (5.3) so that $|\mathcal{N}|$ is as small as possible.

We claim that $(\cup(\mathcal{Z}_i)) \cap (\cup(\mathcal{Z}_j)) = \emptyset$ holds whenever $1 \leq i < j \leq t$. Indeed, $(\cup(\mathcal{Z}_i)) \cap (\cup(\mathcal{Z}_j)) \neq \emptyset$ would imply $|\cup(\mathcal{Z}_i \cup \mathcal{Z}_j)| \leq |\cup(\mathcal{Z}_i)| + |\cup(\mathcal{Z}_j)| - 1$, and hence $b_H(\mathcal{Z}_i \cup \mathcal{Z}_j) \leq b_H(\mathcal{Z}_i) + b_H(\mathcal{Z}_j)$. This is however impossible since by replacing $\mathcal{Z}_i$ and $\mathcal{Z}_j$ in $\mathcal{N}$ by their union we would obtain another minimizer subpartition $\mathcal{N}'$ of $\mathcal{E}$ for which $|\mathcal{N}'| < |\mathcal{N}|$.

We claim furthermore that $e \not\subseteq \cup(\mathcal{Z}_i)$ for every hyperedge $e \in \mathcal{E} - \cup_{i=1}^t \mathcal{Z}_i$. If we had $e \subseteq \cup(\mathcal{Z}_i)$, then by replacing $\mathcal{Z}_i$ by $\mathcal{Z}_i' := \mathcal{Z}_i + e$ we would obtain a subpartition $\mathcal{N}'$ of $\mathcal{E}$ for which $\sum_{\mathcal{Z} \in \mathcal{N}'} b_H(\mathcal{Z}) + |\mathcal{E} - \cup_{\mathcal{Z} \in \mathcal{N}'} \mathcal{Z}| < \sum_{\mathcal{Z} \in \mathcal{N}} b_H(\mathcal{Z}) + |\mathcal{E} - \cup_{\mathcal{Z} \in \mathcal{N}} \mathcal{Z}|$ contradicting the fact that $\mathcal{N}$ is a minimizer.

Let $\mathcal{F}$ be the following partition of V. For each member $\mathcal{Z}_i$ of $\mathcal{N}$, let $\cup(\mathcal{Z}_i)$ be a member of $\mathcal{F}$, and for each element $v \in V - \cup_{i=1}^t \cup (\mathcal{Z}_i)$, let $\{v\}$ be a member of $\mathcal{F}$. By the claims above $e_{\mathcal{E}}(\mathcal{F}) = |\mathcal{E} - \cup_{i=1}^t \mathcal{Z}_i|$. Thus $r_H(\mathcal{E}) = \sum_{i=1}^t b_H(\mathcal{Z}_i) + |\mathcal{E} - \cup_{i=1}^t \mathcal{Z}_i| = \sum_{i=1}^t (|\cup(\mathcal{Z}_i)| - 1) + e_{\mathcal{E}}(\mathcal{F}) = |V| - |\mathcal{F}| + e_{\mathcal{E}}(\mathcal{F})$, as required. □

Corollary 5.5. *The rank of the circuit matroid of a hypergraph $H = (V, \mathcal{E})$ is $|V| - 1$ (i.e. H contains a tree-reducible sub-hypergraph) if and only if H is partition-connected.*

Proof. By definition $r_H(\mathcal{E}) \leq |V| - 1$ and it follows from (5.2) that equality holds precisely if $|V| - |\mathcal{F}| + e_H(\mathcal{F}) \geq |V| - 1$ holds for every partition $\mathcal{F}$ of V, that is, $e_H(\mathcal{F}) \geq |\mathcal{F}| - 1$, which in turn is equivalent to the partition-connectivity of H. □

5.2.3 Decomposition into partition-connected sub-hypergraphs

By Corollary 5.5, a hypergraph can be decomposed into k partition-connected sub-hypergraphs if and only if it contains k edge-disjoint tree-reducible sub-hypergraphs. This can also be described as a matroid problem. For a positive integer k let M_{kH} denote the sum of k copies of the matroid M_H, as defined in Theorem 2.7 of Edmonds. This matroid is defined on the set of hyperedges and a subset of hyperedges is independent if it can be decomposed into k wooded sub-hypergraphs.

Theorem 5.6. *The rank-function r_{kH} of the matroid M_{kH} is given by the following formula:*

$$r_{kH}(\mathcal{X}) = \min\{k(|V| - |\mathcal{F}|) + e_{\mathcal{X}}(\mathcal{F}) : \ \mathcal{F} \textit{ is a partition of } V\}. \tag{5.4}$$

Proof. Again, it suffices to prove the formula for the special case $\mathcal{X} = \mathcal{E}$. Also, as Theorem 5.4 contained the special case $k = 1$, we may assume that $k \geq 2$. As an independent set of

M_{kH} may contain at most $k(|Z|-1)$ hyperedges which are induced by some member Z of $\mathcal{F}$, we see that $r_{kH}(\mathcal{E})$ is at most $k(|V|-|\mathcal{F}|)+e_{\mathcal{E}}(\mathcal{F})$.

To see the reverse inequality we must find a partition $\mathcal{F}$ of V for which $r_{kH}(\mathcal{E}) = k(|V|-|\mathcal{F}|)+e_{\mathcal{E}}(\mathcal{F})$. Equation (2.3) in Theorem 2.7 gives the rank function of the sum of matroids. Combining this and (5.2) we obtain that $r_{kH}(\mathcal{E}) = \min_{\mathcal{E}'\subseteq\mathcal{E}}(kr_H(\mathcal{E}')+|\mathcal{E}-\mathcal{E}'|) = \min_{\mathcal{E}'\subseteq\mathcal{E}}(k(\min\{|V|-|\mathcal{F}|+e_{\mathcal{E}'}(\mathcal{F}) : \mathcal{F} \text{ is a partition of } V\}) \ +|\mathcal{E}-\mathcal{E}'|]$. Let $\mathcal{E}'$ be a subset of $\mathcal{E}$ for which the minimum is attained and let $\mathcal{F}$ be a partition of V for which the inner minimum is attained.

We claim that $e_{\mathcal{E}'}(\mathcal{F}) = 0$. Indeed, if there was an element $e \in \mathcal{E}'$ with neighbours in at least two members of $\mathcal{F}$, then for $\mathcal{E}'' := \mathcal{E}' - e$ we would have $ke_{\mathcal{E}''}(\mathcal{F}) + |\mathcal{E}-\mathcal{E}''| = ke_{E'}(\mathcal{F}) + |\mathcal{E}-\mathcal{E}'| - (k-1)$ and, since $k \geq 2$, this would contradict the assumption that $\mathcal{E}'$ is a minimizer.

We also claim that every edge $e \in \mathcal{E}-\mathcal{E}'$ is a cross-edge of $\mathcal{F}$, for otherwise $e_{\mathcal{E}''}(\mathcal{F}) = e_{\mathcal{E}'}(\mathcal{F})$ would hold for $\mathcal{E}'' := \mathcal{E}' + e$ for some $e \in \mathcal{E}-\mathcal{E}'$, which is again impossible if $\mathcal{E}'$ is a minimizer.

We can conclude that $\mathcal{E}-\mathcal{E}'$ consists of exactly the cross-edges of $\mathcal{F}$ in $\mathcal{E}$. Hence $r_{kH}(\mathcal{E}) = k(|V|-|\mathcal{F}|+e_{\mathcal{E}'}(\mathcal{F})) + |\mathcal{E}-\mathcal{E}'| = k(|V|-|\mathcal{F}|)+e_{\mathcal{E}}(\mathcal{F})$, as required. □

Now we can establish the link between k-partition-connectivity and the decomposability into partition-connected sub-hypergraphs. The following result generalizes Theorem 1.5 of Tutte.

Theorem 5.7. *For a hypergraph $H = (V, \mathcal{E})$ the following are equivalent:*

(i) H is k-partition-connected,

(ii) H can be decomposed into k partition-connected sub-hypergraphs,

(iii) H contains k edge-disjoint tree-reducible spanning sub-hypergraphs.

Proof. If the decomposition into k partition-connected sub-hypergraphs exists, the hypergraph is clearly k-partition-connected. To see the other direction, suppose that H is k-partition-connected, that is, $e_{\mathcal{E}}(\mathcal{F}) \geq k(|\mathcal{F}|-1)$ holds for every partition $\mathcal{F}$ of V. This is equivalent to requiring that $k(|V|-|\mathcal{F}|)+e_{\mathcal{E}}(\mathcal{F}) \geq k(|V|-1)$, i.e. $r_{kH} = k(|V|-1)$ by Theorem 5.6. Therefore every basis of the matroid M_{kH} is the union of k bases of M_H of cardinality $|V|-1$. It follows that H contains k edge-disjoint tree-reducible spanning sub-hypergraphs, which by Corollary 5.5 is equivalent to the property that H can be decomposed into k partition-connected spanning sub-hypergraphs. □

A well-known corollary of Theorem 1.5 of Tutte is that a $(2k)$-edge-connected graph always contains k disjoint spanning trees. As a direct extension of this result we can derive the following:

Corollary 5.8. *A (νk)-edge-connected hypergraph H of rank at most ν can be decomposed into k partition-connected sub-hypergraphs and hence into k connected spanning sub-hypergraphs.*

Proof. By Theorem 5.7 it suffices to show that H is k-partition-connected. Let $\mathcal{F}$ be a partition of V. By the (νk)-edge-connectivity of H, there are at least νk hyperedges intersecting both X and $V-X$ for each member X of $\mathcal{F}$. Since one hyperedge may intersect at most ν members, we obtain that the total number of hyperedges intersecting more than one member of $\mathcal{F}$ is at least $(\nu k)|\mathcal{F}|/\nu \geq k(|\mathcal{F}|-1)$, so H is k-partition-connected. □

Note that the proof shows that a bit stronger result is also true: if H is a (νk)-edge-connected hypergraph of rank at most ν, then after the deletion of any k hyperedges, H still decomposes into k partition-connected sub-hypergraphs. This kind of connectivity property is similar to the (k,l)-partition-connectivity of graphs introduced in Chapter 1, and indeed it is useful to extend this notion to hypergraphs.

5.2.4 (k,l)-partition-connectivity

For non-negative integers k and l, a hypergraph $H=(V,\mathcal{E})$ is called (k,l)-*partition-connected* if $e_H(\mathcal{F}) \geq k(|\mathcal{F}|-1)+l$ for every nontrivial partition $\mathcal{F}$. An equivalent definition is that no matter how we delete l hyperedges from H, the resulting hypergraph decomposes into k partition-connected sub-hypergraphs. This may be regarded as a combination of two different approaches to connectivity: that the hypergraph should be decomposable into many well-connected components, and that the hypergraph should remain well-connected even after the deletion of some hyperedges.

A notable difference from the graph case is that while (k,l)-partition-connectivity for $k \leq l$ is equivalent to $(k+l)$-edge-connectivity for graphs, the same is not true for hypergraphs. As a consequence, the case $k<l$ also has some interest.

Note that if $|\mathcal{E}|=|V|$ for a hypergraph $H=(V,\mathcal{E})$, then H is $(1,1)$-partition-connected if and only if it is a hypercircuit. Indeed, suppose that $i_H(X) \geq |X|$ for some $X \subset V$. Then for the partition $\mathcal{F} := \{X\} \cup \{\{v\} : v \in V-X\}$ we have $e_H(\mathcal{F}) \leq |\mathcal{E}|-i_H(X) \leq |V|-|X| \leq |\mathcal{F}|-1$. Conversely, if there is a partition for which $e_H(\mathcal{F}) \leq |\mathcal{F}|-1$, then by assuming that the strong Hall condition holds we get that $|V|=|\mathcal{E}|=e_H(\mathcal{F})+\sum_{Z\in\mathcal{F}} i_H(Z) \leq |\mathcal{F}|-1+\sum_{Z\in\mathcal{F}}(|Z|-1) \leq |V|-1$, a contradiction.

There is a difference in the structural properties of k-partition-connected hypergraphs, and (k,l)-partition-connected hypergraphs in general. Namely, Theorem 5.7 implies that a k-partition-connected hypergraph is always reducible to a k-partition-connected graph, in the sense used in the definition of forest-reducibility. However, a (k,l)-partition-connected hypergraph is not necessarily reducible to a (k,l)-partition-connected graph. For example, the hypercircuits on Figure 5.2 are not reducible (of course there are also reducible hypercircuits).

If we want to prove someone that a given hypergraph is k-partition-connected, we can do so by showing its reduction to a k-partition-connected graph, which contains k edge-disjoint spanning trees. But how can we prove that a given hypergraph is (k,l)-partition-connected? An answer to this question will be given in Chapter 6.

5.2.5 Covering by hyperforests

The problems discussed so far in this chapter were generalizations of problems that involve the packing of trees. It may be asked how problems related to covering with forests generalize to hypergraphs. An extension of Theorem 1.6 of Nash-Williams on partitioning a graph into k forests easily follows from Theorem 5.6.

Theorem 5.9. *The edge-set $\mathcal{E}$ of a hypergraph $H = (V,\mathcal{E})$ can be decomposed into k hyperforests if and only if*

$$i_H(X) \le k(|X|-1) \tag{5.5}$$

holds for every nonempty subset X of V.

Proof. Since one hyperforest may contain at most $|X|-1$ hyperedges induced by X, (5.5) is necessary. To see the sufficiency, let $\mathcal{F} := \{V_1,\dots,V_t\}$ be a partition of V. By (5.5) each member V_i induces at most $k(|V_i|-1)$ hyperedges, so the number $e_H(\mathcal{F})$ of cross-hyperedges is at least $|\mathcal{E}| - \sum_{j=1}^{t} k(|V_j|-1) = |\mathcal{E}| - k(|V|-t)$. Therefore $k(|V|-|\mathcal{F}|) + e_H(\mathcal{F}) \ge |\mathcal{E}|$ holds for every partition $\mathcal{F}$ of V, which means that $\mathcal{E}$ is independent in the matroid M_{kH} by (5.4), and H decomposes into k hyperforests. □

The problem of connectivity augmentation with respect to (k,l)-partition-connectivity will be discussed in Chapter 7. Here we address the much easier analogous problem concerning covering with hyperforests: add hyperedges to a given hypergraph $H_0 = (V,\mathcal{E}_0)$ such that the resulting hypergraph can still be decomposed into k hyperforests. It is a simple observation that even the maximum weight problem is solvable (assuming that the number of hyperedges with positive weight is polynomial), by the matroid techniques described in the previous sections. Indeed, let us assign a sufficiently large weight N to every $e \in \mathcal{E}_0$, and let $\mathcal{E}^*$ be the set of hyperedges with positive weight, and $H^* = (V, \mathcal{E}^*)$. Then finding a maximum weight feasible augmentation amounts to finding a maximum weight independent set in the matroid M_{kH^*}.

The following theorem gives a min-max formula on the corresponding degree-specified augmentation problem, when we allow only the addition of graph edges:

Theorem 5.10. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $m : V \to \mathbb{Z}_+$ a degree specification with $m(V)$ even, and k a positive integer. There exists an undirected graph $G = (V, E)$ such that $H_0 + G$ can be covered by k hyperforests and $d_G(v) = m(v)$ for every $v \in V$ if and only if*

$$\left(m(X) - \frac{m(V)}{2}\right)^+ \leq k(|X| - 1) - i_{H_0}(X) \text{ for every } \emptyset \neq X \subseteq V. \tag{5.6}$$

Proof. Clearly, $i_G(X) \geq (m(X) - \frac{m(V)}{2})^+$, so the requirement $i_{H_0+G}(X) \leq k(|X| - 1)$ implies the necessity of condition (5.6).

We prove sufficiency by induction on $m(V)$. By Theorem 5.9, a hypergraph can be covered by k hyperforests if and only if $i_{H_0}(X) \leq k(|X| - 1)$ for every non-empty subset X of V; hence we may assume that $m(V) \geq 2$. Let $v \in V$ be an arbitrary node with $m(v) > 0$.

A set X is called *tight* if (5.6) holds with equality. Let $\mathcal{F}_1$ be the family that consists of the tight sets X for which $m(X) \leq m(V)/2$ and $v \in X$, and let $\mathcal{F}_2$ be the family of tight sets X for which $m(X) \geq m(V)/2$ and $v \notin X$. The union of two members of $\mathcal{F}_1$ is also in $\mathcal{F}_1$, since otherwise by the supermodularity of i_{H_0} the intersection would violate (5.6); similarly, the intersection of two sets in $\mathcal{F}_2$ is in $\mathcal{F}_2$, since otherwise by the supermodularity of i_{H_0} their union would violate (5.6). Let X_1 be the maximal member of $\mathcal{F}_1$, and X_2 the minimal member of $\mathcal{F}_2$. Then $v \in X_1 - X_2$ and $m(X_1) \leq m(X_2)$, so there is a node $u \in X_2 - X_1$ with $m(u) > 0$.

Let m' be defined by decreasing $m(u)$ and $m(v)$ by 1, and H_0' defined by adding the graph edge uv to H_0. The node u was chosen so that no member of $\mathcal{F}_1$ contains both u and v, and every member of $\mathcal{F}_2$ contains u. From this it is easy to see that (5.6) holds for m' and H_0', therefore H_0' can be augmented by adding a graph G' with degree-specification m' such that $H_0' + G'$ can be covered by k forests. This means that $G := G' + \{uv\}$ is a good augmenting graph for H_0. ☐

5.3 Disjoint Steiner trees

Let $G = (V, E)$ be an undirected graph with a so-called terminal set $W \subseteq V$. In Chapter 1, we defined a *Steiner tree* for W as a subtree $G' = (V', E')$ of G such that $W \subseteq V'$. The *disjoint Steiner trees problem* consists of finding k edge-disjoint Steiner trees in G. When $W = V$, this corresponds to the existence of k disjoint spanning trees and Theorem 1.5 provides a characterization. When $|W| = 2$, a minimal Steiner tree is a path connecting the two terminal nodes, and Theorem 1.1 gives an answer. However, as it was mentioned in Chapter 1, for general W it is NP-complete to decide whether G contains k edge-disjoint Steiner trees for W.

This means that deriving sufficient conditions for the existence of k disjoint Steiner trees may be of some interest. One type of possible sufficient condition is high edge-connectivity in W. Kriesell [53] proposed the conjecture that $2k$-edge-connectivity in W is sufficient for the existence of k edge-disjoint Steiner trees. The conjecture is open even for an arbitrary constant multiple of k instead of 2. Jain, Mahdian, and Salavatipour proved the following:

Theorem 5.11 (Jain et al. [46]). *If G is l-edge-connected in W then G contains $\lfloor \alpha_{|W|} l \rfloor$ edge-disjoint Steiner trees for W, where α_i can be defined recursively by*

$$\alpha_2 = 1,$$
$$\alpha_i = \alpha_{i-1} - \frac{\alpha_{i-1}^2}{4} \quad \text{for } i > 2.$$

In the following we show using Corollary 5.8 that if $V - W$ is stable, then $3k$-edge-connectivity in W of G is sufficient for the existence of k edge-disjoint Steiner trees.

Theorem 5.12. *Let $G = (V, E)$ be an undirected graph and $W \subset V$ a subset of nodes so that $U := V - W$ is stable and G is $(3k)$-edge-connected in W. Then G contains k edge-disjoint Steiner trees for W.*

Proof. We use induction on the value $\mu_G := \sum_{u \in U} (d_G(u) - 3)^+$. Suppose first that μ_G is zero, that is, the degree of each node in U is at most 3. We may assume that W is also stable for otherwise each edge induced by W can be subdivided by a new node. Such an operation may add new nodes of degree two to U but it does not affect $(3k)$-edge-connectivity in W and k disjoint Steiner trees in the new graph determine k disjoint Steiner trees in G.

Let $H = (W, \mathcal{E})$ be the hypergraph corresponding to the bipartite graph $G = (U, W; E)$, i.e. for each node $u \in U$ there is a corresponding hyperedge of H consisting of the neighbours of u in G. As the degree of each node of U is at most 3, the rank of H is at most 3. For any $\emptyset \neq X \subset W$ let X' denote the set of those nodes of U which have at least one neighbour in X and at most one neighbour in $W - X$ in the graph G. Since every degree in U is at most 3, we have $d_G(X \cup X') = d_H(X)$ and hence the $(3k)$-edge-connectivity of G in W implies the $(3k)$-edge-connectivity of H.

By Corollary 5.8, H can be decomposed into k connected spanning sub-hypergraphs, thus U can be partitioned into k disjoint subsets $U_1, \ldots, U_k$ so that $W \cup U_i$ induces a connected subgraph $G_i = (W \cup U_i, E_i)$ of G for each $i = 1, \ldots, k$. By choosing one spanning tree from each G_i, we obtain the required edge-disjoint Steiner trees of G.

Suppose now that μ_G is positive and that the theorem holds for each graph G' with $\mu_{G'} < \mu_G$. Let $s \in U$ be a node with $d_G(s) \geq 4$. If there is a cut-edge e of G, then the elements of W belong to the same component of $G - e$ as G is at least $(3k)$-edge-connected in W and then we may discard the other component of $G - e$ without destroying $(3k)$-edge-connectivity in W. Therefore we may assume that G is 2-edge-connected.

By Theorem 3.3 of Mader on splitting-off preserving local edge-connectivity, there are two edges $e_1 = v_1 s, e_2 = v_2 s$ in E so that the local edge-connectivities in $V - s$ do not decrease if we replace e_1 and e_2 by a new edge $v_1 v_2$. In particular, the resulting graph G' remains $(3k)$-edge-connected in W. By induction there are k edge-disjoint Steiner trees for W in G'. If one of these trees contains the split-off edge $v_1 v_2$, we replace it by e_1 and/or e_2 in order to obtain a Steiner tree of G. Therefore we have proved the existence of k edge-disjoint Steiner trees for W in G. □

It should be noted that if the degree of each node in U is even, then $2k$-edge-connectivity in W implies the existence of k edge-disjoint Steiner trees for W, even without the restriction that U must be stable. This can be seen by observing that in this case Theorem 3.3 of Mader can always be applied since we can assume that G is 2-edge-connected, and the splitting-off preserves the parity of the degrees, therefore the degree of a node in U will never be 3.

As far as algorithmic aspects are concerned, Edmonds' matroid partition algorithm may be used to compute a decomposition of a hypergraph into k partition-connected sub-hypergraphs or to compute a deficient partition to show that such a decomposition does not exist. Therefore Theorem 5.12 can be used for an approximation algorithm to compute the maximum number τ of disjoint Steiner trees when $V - W$ is stable. W is clearly τ-edge-connected in G, so by Theorem 5.12 we can compute $\tau/3$ disjoint Steiner trees.

Chapter 6

Hypergraph orientation

6.1 Introduction

In Chapter 4 we saw that many results on the edge-connectivity of digraphs are extendable to directed hypergraphs. It is a natural question to ask whether similar extensions are possible for connectivity orientation problems. For example, when does a hypergraph have an orientation that is (k,l)-edge-connected? This Chapter is devoted to questions of this type, and we also present some orientation results that are new even for graphs.

In Section 6.2 we briefly review known connectivity orientation results for graphs. In particular, we describe the relation between high partition-connectivity of a graph and the existence of highly edge-connected orientations. One of the observations in Section 6.3 is that a similar relation exists in the case of hypergraphs. The reason behind the similarities is the link between orientation problems and submodularity, which can be established for hypergraphs just like for graphs.

In [50], Khanna, Naor and Shepherd proposed a new framework, called network design with orientation constraints, that successfully integrated network design problems like minimum cost k-rooted-connected sub-digraphs, and orientation problems like k-rooted-connected orientation of a mixed graph. In Section 6.4 we extend their result to hypergraphs, and show that their formulation actually defines a TDI system, which gives rise to new min-max formulas.

Orientations with local edge-connectivity properties define an exciting class of problems. For graphs we have the beautiful theorem of Nash-Williams [61] (Theorem 6.4) on well-balanced orientations; a major open question of the field is how to generalize this theorem to hypergraphs. The new results in Section 6.5 are rather simple and do not really mean a step forward in that direction; however, they do have some nice corollaries, including a new characterization of $(2k+1)$-edge-connected graphs, and characterization of (k,l)-

partition-connected hypergraphs if $k < l$.

The new results in this chapter are based on [36], a joint work with András Frank and Zoltán Király.

6.2 Orientation of graphs

6.2.1 Connectivity requirements

In a graph connectivity orientation problem one is interested in the existence of an orientation of an undirected graph that satisfies some specified connectivity requirements. One of the early examples is the theorem of Robbins [65]: a graph has a strongly connected orientation if and only if it is 2-edge-connected. As an illustration, he rephrased this as a traffic control problem: decide whether the streets of a city can be turned into one-way streets such that any location in the city remains reachable from any other location. Nash-Williams [61] proved the following generalization of this result:

Theorem 6.1 (Nash-Williams [61], weak form). *A graph has a k-edge-connected orientation if and only if it is $2k$-edge-connected.*

Another type of connectivity orientation result follows easily from Theorem 1.5 of Tutte: as we have seen in Chapter 1, a graph has a k-rooted-connected orientation (from an arbitrary root s) if and only if it is k-partition-connected.

It is natural to try to extend these results to (k, l)-edge-connectivity, since (k, k)-edge-connectivity of a digraph is equivalent to k-edge-connectivity, and $(k, 0)$-edge-connectivity is equivalent to k-rooted-connectivity from some node s (note that if a graph has a (k, l)-edge-connected orientation from some root, then it has such an orientation from any root). This problem was solved in [22]:

Theorem 6.2 (Frank [22]). *Let $k \geq l$ be non-negative integers. A graph has a (k, l)-edge-connected orientation if and only if it is (k, l)-partition-connected.*

Actually this was proved in a more general framework, involving orientations that cover a given non-negative crossing supermodular set function:

Theorem 6.3 (Frank [22]). *Let $G = (V, E)$ be a graph, and $p : 2^V \to \mathbb{Z}_+$ a non-negative crossing supermodular set function. Then G has an orientation covering p if and only if*

$$\sum_{X \in \mathcal{F}} p(X) \leq e_G(\mathcal{F}), \tag{6.1}$$

$$\sum_{X \in \mathcal{F}} p(V - X) \leq e_G(\mathcal{F}) \tag{6.2}$$

both hold for every partition $\mathcal{F}$ of V. If p is monotone decreasing or symmetric then it suffices to require the validity of (6.1).

The necessity of these conditions is obvious if we recall (1.7), and that $e_G(\mathcal{F}) = e_G(\text{co}(\mathcal{F}))$ if G is a graph (but this is not true for hypergraphs!). The result on (k,l)-edge-connected orientations follows from this theorem by considering the following set function for an arbitrary node $s \in V$:

$$p_{kl}(X) := \begin{cases} k & \text{if } \emptyset \neq X \subseteq V - s, \\ l & \text{if } s \in X \subset V, \\ 0 & \text{otherwise.} \end{cases} \tag{6.3}$$

It is easy to see that p_{kl} is non-negative, monotone decreasing if $k \geq l$, and crossing supermodular.

For orientations satisfying local edge-connectivity requirements, the problem is much more difficult. Let $G = (V,E)$ be a graph, and let $r : V^2 \to \mathbb{Z}_+$ be a local edge-connectivity requirement function for which $r(v,v) = 0$ $(v \in V)$. If r is symmetric, i.e. $r(u,v) = r(v,u)$ for every $u,v \in V$, then an obvious necessary condition for the existence of a good orientation is that

$$r(u,v) \leq \left\lfloor \frac{\lambda_G(u,v)}{2} \right\rfloor \quad \text{for every } u,v \in V.$$

The orientation theorem of Nash-Williams [61] states that this condition is sufficient for symmetric r:

Theorem 6.4 (Nash-Williams [61], strong form). *Let $G = (V,E)$ be an arbitrary graph. There exists an orientation $\vec{G} = (V,\vec{E})$ of G such that*

$$\lambda_{\vec{G}}(u,v) \geq \left\lfloor \frac{\lambda_G(u,v)}{2} \right\rfloor \quad \textit{for every } u,v \in V,$$

and

$$|\varrho_{\vec{G}}(v) - \delta_{\vec{G}}(v)| \leq 1 \quad \textit{for every } v \in V.$$

This theorem settles the case when r is symmetric. However, deciding whether there is an orientation satisfying a general local edge-connectivity requirement r is NP-complete. The following is a sketch of the reduction of 3-SAT (see Figure 6.1).

Consider a collection $\mathcal{C}$ of clauses, and construct the following graph G. For every pair $\{x,\overline{x}\}$ of complementary literals, create two nodes v_x and $v_{\overline{x}}$, and an edge $v_x v_{\overline{x}}$. For each clause $c \in \mathcal{C}$, add nodes s_c, t_c, w_c, z_c; for each literal $y \in c$, add edges $v_{\overline{y}} s_c$, $v_y t_c$, $v_{\overline{y}} w_c$, and

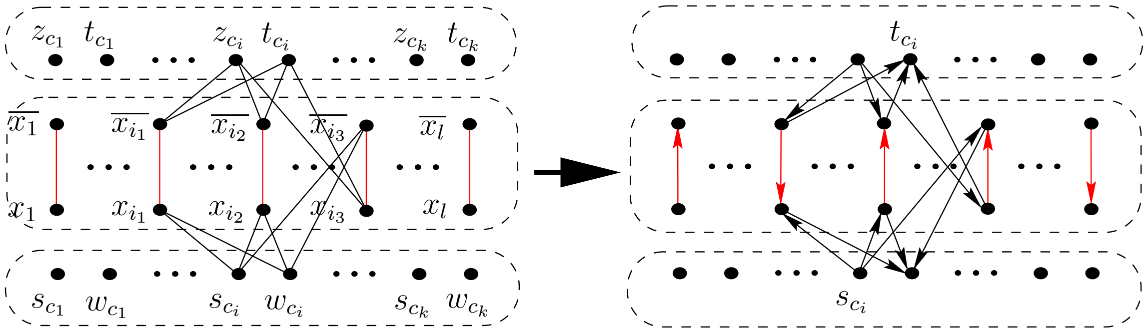

Figure 6.1: Reduction of 3-SAT to local edge-connectivity orientation. The figure shows the construction for the clause $c_i = (\overline{x_{i_1}}, \overline{x_{i_2}}, x_{i_3})$. An orientation of the edges of type $x\overline{x}$ corresponds to an evaluation; the orientation of the other edges is uniquely determined.

$v_y z_c$. Consider the problem of finding an orientation of G such that for every clause $c \in \mathcal{C}$ there are at least 3 edge-disjoint paths from s_c to w_c, 3 edge-disjoint paths from z_c to t_c, and 1 path from s_c to t_c. It is easy to see that the existence of such an orientation is equivalent to the satisfiability of $\mathcal{C}$.

6.2.2 Orientations and submodular flows

The orientation problems described above can be studied for mixed graphs (graphs with both undirected and directed edges) as well. An orientation of a mixed graph is obtained by orienting its undirected edges. In [25] Frank solved the problem of finding an orientation of a mixed graph that covers a given crossing supermodular set function which does not have to be non-negative. He showed that this is equivalent to a submodular flow problem, and as a result minimum cost orientation (when the two possible orientations of an undirected edge have different costs) can be solved in polynomial time. The characterization here is considerably more complicated then in Theorem 6.3:

Theorem 6.5 (Frank [25]). *Let $G = (V; E, A)$ be a mixed graph, where E is the set of undirected edges, and A is the set of directed edges. Let $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ be a crossing supermodular set function. Then G has an orientation covering p if and only if*

$$\sum_{Z \in \mathcal{F}} (p(Z) - \varrho_A(Z)) \leq e_E(\mathcal{F})$$

holds whenever $\mathcal{F}$ is a tree-composition of some $X \subseteq V$.

The reason that tree-compositions come into the picture is the connection with submodular flows: in Theorem 2.28 we saw that the condition on the feasibility of a submodular flow problem given by a crossing supermodular set function involves the full truncation of the set function.

6.2.3 Constructive characterizations

One important application of the connection between the connectivity properties of undirected graphs and the connectivity properties of their orientations concerns the theory of constructive characterizations. Constructive characterization of a connectivity property means that such (di)graphs can be constructed from a small initial (di)graph by some simple operations that preserve the property.

For (k,l)-edge-connectivity of digraphs, the following conjecture is expected to be true:

Conjecture 6.6. *Let $k > l$ be non-negative integers. A digraph is (k,l)-edge-connected if and only if it can be built from a single node by the following two operations:*

(i) Add a new edge to the digraph,

(ii) Pinch i ($l \leq i < k$) existing edges with a new node z, and add $k-i$ new edges entering z and leaving existing nodes,

where pinching edges $u_1v_1, \dots, u_tv_t$ with z means deleting those edges, and adding the edges $u_1z, zv_1, u_2z, zv_2, \dots, u_tz, zv_t$.

It is easy to see that the above operations create a (k,l)-edge-connected digraph. The conjecture was proved for $l = 1$ by Frank and Szegő [34], and for $l = k-1$ by Frank and Z. Király [31]. It is also easy to see that since for $k > l$ (k,l)-partition-connected graphs are exactly those that have a (k,l)-edge-connected orientation, proof of this conjecture would also imply the following on the constructive characterization of (k,l)-partition-connected graphs:

Conjecture 6.7. *Let $k > l$ be non-negative integers. A graph is (k,l)-partition-connected if and only if it can be built from a single node by the following two operations:*

(i) Add a new edge to the graph,

(ii) Pinch i ($l \leq i < k$) existing edges with a new node z, and add $k-i$ new edges connecting z with existing nodes.

6.2.4 Parity constrained orientation

A very exciting new class of connectivity orientation problems is obtained if there are additional requirements involving the parity of the in-degrees in the oriented graph. These problems are outside the scope of this thesis, so we just mention a few results.

Let $G = (V, E)$ be a graph. For a subset $T \subseteq V$, a *T-odd orientation* of G is an orientation where T is exactly the set of nodes with odd in-degrees. It is easy to see that a T-odd orientation can exist only if $|T| + |E|$ is even; such sets are called *G-even.* Frank, Jordán and Szigeti proved the following on k-rooted-connected T-odd orientations:

Theorem 6.8 (Frank et al. [33]). *Let $G = (V, E)$ be a graph with a fixed root node $s \in V$, let k be a positive integer, and let $T \subseteq V$ be a G-even subset. For a partition $\mathcal{F}$ of V, let $\mathrm{odd}(\mathcal{F})$ denote the number of members of $\mathcal{F}$ for which $|X \cap T| - i_G(X) - k$ is odd. Then G has a k-rooted-connected T-odd orientation if and only if*

$$k(|\mathcal{F}| - 1) + \mathrm{odd}(\mathcal{F}) \le e_G(\mathcal{F})$$

for every partition $\mathcal{F}$ of V with $|\mathcal{F}| \ge 2$.

It is open (even for $k = 1$) whether graphs with a k-edge-connected T-odd orientation can be characterized. However, Frank and Z. Király [31] managed to characterize graphs that have a k-edge-connected T-odd orientation for every G-even subset T.

Theorem 6.9 (Frank, Z. Király [31]). *A graph $G = (V, E)$ has a k-edge-connected T-odd orientation for every G-even subset T if and only if it is $(k+1, k)$-partition-connected.*

The proof of this theorem is a very nice application of constructive characterizations.

6.3 Orientation with supermodular requirements

Now we turn to orientation problems featuring hypergraphs. This section presents hypergraph analogues of Theorems 6.2 and 6.3. We start by describing the framework used for hypergraph connectivity orientation.

6.3.1 General framework

Hypergraph orientation problems can be formulated in basically the same way as their graph counterparts. Let $H = (V, \mathcal{E})$ be a hypergraph, and $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ a set function, called the *requirement function.* We consider the problem of finding an orientation of H that covers p. For example, if p is the set function p_{kl} defined in (6.3), then this amounts to finding a (k,l)-edge-connected orientation of H (the notion of (k,l)-edge-connectivity of directed hypergraphs was defined in Chapter 4). Any local edge-connectivity requirement can also be formulated in this way.

One difference from the graph case is that in a directed hypergraph the role of the head node and that of the tail nodes are not symmetric, therefore an orientation cannot be reversed. It should also be mentioned that if we consider the bipartite graph associated to a hypergraph as in Chapter 5, then hypergraph orientation can be seen as a kind of constrained orientation problem for bipartite graphs, where every node in one class should have out-degree exactly 1. So these problems could be described in the graph framework; however, the hypergraph formulation is much more comfortable.

6.3.2 Orientation Lemma

The easiest orientation problem is when there is no connectivity requirement, but the in-degree of every node is specified. The following hypergraph orientation lemma, which characterizes the existence of such an orientation, is a straightforward extension of the corresponding graph result (see e.g. [44]). Note that this is not a real generalization, since it follows from the graph case by considering the bipartite graph associated to the hypergraph. Here we give a direct proof.

Lemma 6.10. *Given a hypergraph H and an in-degree specification $m_i : V \to \mathbb{Z}_+$, there is an orientation $\vec{H}$ of H such that $\varrho_{\vec{H}}(v) = m_i(v)$ for every $v \in V$ if and only if $m_i(V) = |\mathcal{E}|$ and $m_i(Y) \geq i_H(Y)$ for every $Y \subseteq V$.*

Proof. The necessity is straightforward. We prove the sufficiency by induction on the number of hyperedges. Call a set Y *tight* if $m_i(Y) = i_H(Y)$. Let $e \in \mathcal{E}$ be an arbitrary hyperedge, and let Z be the set of nodes in e. Then $m_i(Z - X) \geq 1$ for any tight set X for which $Z \not\subseteq X$ (including $X = \emptyset$), otherwise $Z \cup X$ would violate the condition. If there is a node $v \in Z$ with $m_i(v) > 0$ such that $Z \subseteq X$ for every tight set X containing v, then we can remove the hyperedge e, decrease $m_i(v)$ by one, find a feasible orientation of the resulting hypergraph by induction, and add the oriented hyperedge $(e, h(e))$ with $h(e) = v$. Otherwise, since a single tight set $X \not\supseteq Z$ cannot contain every node $v \in Z$ with $m_i(v) > 0$, we can choose tight sets X_1 and X_2 that are both maximal among the tight sets that separate Z. Then $X_1 \cup X_2$ is tight and $d_H(X_1, X_2) = 0$ because of (1.3), therefore $Z - (X_1 \cup X_2) \neq \emptyset$, which contradicts the maximality of X_1 and X_2. $\square$

To see the importance of this Lemma, observe that the in-degrees of an orientation determine whether the orientation covers a given set function p or not. More precisely, an orientation of H with in-degree specification m_i covers p if and only if $m_i(X) \geq i_H(X) + p(X)$ for every $X \subseteq V$.

6.3.3 Covering crossing supermodular set functions

We are now ready to solve the hypergraph orientation problem when the requirement function p is non-negative and crossing supermodular.

Theorem 6.11. *Let $H = (V, \mathcal{E})$ be a hypergraph, and $p : 2^V \to \mathbb{Z}_+$ a non-negative crossing supermodular set function. There is an orientation of H covering p if and only if*

$$\sum_{X \in \mathcal{F}} p(X) \leq e_H(\mathcal{F}) \tag{6.4}$$

holds for every partition and co-partition $\mathcal{F}$.

Proof. The necessity of the conditions can be seen by considering property (1.7) of e_H. It is important to realize that here $e_H(\mathcal{F})$ and $e_H(\text{co}(\mathcal{F}))$ can be different. If $\mathcal{F}$ is a partition of V, then $e_H(\mathcal{F}) \leq e_H(\text{co}(\mathcal{F}))$.

We turn to the proof of sufficiency. We may assume that every hyperedge of H contains at least 2 distinct nodes. A partition or a co-partition $\mathcal{F}$ is called *tight* if $\sum_{X \in \mathcal{F}} p(X) = e_H(\mathcal{F})$. Observe that the crossing supermodularity remains valid if we increase the value of p on some singletons; we can thus assume that every singleton $\{v\}$ is member of a tight partition $\mathcal{F}_v$. Let $\mathcal{F} := \sum_{v \in V} \mathcal{F}_v$ be the sum of these tight partitions, which is a regular family; then $\sum_{X \in \mathcal{F}} p(X) = e_H(\mathcal{F})$. Our aim is to show that this implies $\sum_{v \in V} p(\{v\}) = |\mathcal{E}|$. We can uncross $\mathcal{F}$ using the standard uncrossing operation (see Lemma 2.4), to obtain a cross-free regular family $\mathcal{F}'$ including all the singletons, for which

$$\sum_{X \in \mathcal{F}'} p(X) \geq e_H(\mathcal{F}'), \tag{6.5}$$

since an uncrossing step does not decrease $\sum_{X \in \mathcal{F}} p(X)$ due to the crossing supermodularity of p, and it does not increase $e_H(\mathcal{F})$.

Let $\mathcal{F}''$ be the family obtained by decreasing the multiplicity of every singleton in $\mathcal{F}'$ by 1. By Proposition 2.1, $\mathcal{F}''$ decomposes into partitions and co-partitions, and by Proposition 1.17, (6.4), and (6.5), these must be tight partitions and co-partitions, and the partition formed of singletons is tight as well. As a consequence, if we define $x(v) := p(\{v\})$ for every $v \in V$, then $x(V) = |\mathcal{E}|$.

To complete the proof, it suffices to show that $x(Y) \geq i_H(Y) + p(Y)$ for every set $Y \subseteq V$, since in this case by Lemma 6.10 and the non-negativity of p there is an orientation with in-degree vector x, and this orientation covers p as every set Y is entered by $x(Y) - i_H(Y)$

hyperarcs. To prove the inequality, define the partition $\mathcal{F}_Y := \{Y\} \cup \{\{v\} :\ v \in V - Y\}$ for every set $Y \subset V$. Using (6.4) on the partition $\mathcal{F}_Y$, we get

$$\begin{aligned} x(Y) &= |\mathcal{E}| - x(V - Y) = |\mathcal{E}| + p(Y) - \sum_{Z \in \mathcal{F}_Y} p(Z) \\ &\geq |\mathcal{E}| - e_H(\mathcal{F}_Y) + p(Y) = i_H(Y) + p(Y). \end{aligned}$$

□

It should be noted that using Theorem 2.11 of Fujishige, a short alternative proof of Theorem 6.11 can be given. Define the set function $q(X) := p(X) + i_H(X)$; then q is crossing supermodular. If $\mathcal{F} = \{X_1, \dots, X_t\}$ is a partition of V, then, by (6.4) and (1.12), $\sum q(X_i) = \sum p(X_i) + \sum i_H(X_i) \leq e_H(\mathcal{F}) + \sum i_H(X_i) = |\mathcal{E}| = q(V)$, and $\sum q(V - X_i) = \sum p(V - X_i) + \sum i_H(V - X_i) \leq e_H(\text{co}(\mathcal{F})) + \sum i_H(V - X_i) = (t-1)|\mathcal{E}| = (t-1)q(V)$. Thus Theorem 2.11 implies that if the conditions (6.4) hold, then there is an integral vector $x : V \to \mathbb{Z}$ satisfying $x(V) = q(V) = |\mathcal{E}|$ and $x(Y) \geq q(Y) = i_H(Y) + p(Y) \geq i_H(Y)\ \forall Y \subseteq V$. By Lemma 6.10, H has an orientation with in-degree vector x, and this orientation covers p.

Remark. The proofs show that Theorem 6.11 is true under the weaker assumption that p is non-negative and $p + i_H$ is crossing supermodular. If p is monotone decreasing or symmetric, then the co-partition type constraints are unnecessary, since $\sum_{X \in \mathcal{F}} p(X) \geq \sum_{X \in \text{co}(\mathcal{F})} p(X)$ and $e_H(\mathcal{F}) \leq e_H(\text{co}(\mathcal{F}))$ for every partition $\mathcal{F}$.

6.3.4 Characterization of (k, l)-partition-connectivity $(k \geq l)$

By applying Theorem 6.11 to the set function p_{kl} defined in (6.3), we get a new characterization of (k, l)-partition-connected hypergraphs (as defined in Chapter 5) if $k \geq l$. In this case the set function p_{kl} is monotone decreasing, so by taking into account the previous remark, we obtain the following:

Corollary 6.12. *Let $k \geq l$ be non-negative integers. Then a hypergraph has a (k, l)-edge-connected orientation if and only if it is (k, l)-partition-connected.*

Note that in the special case when $l = 0$, this gives an alternative proof for Theorem 5.7, since by Proposition 1.14 a k-rooted-connected orientation can be decomposed into k rooted-connected directed sub-hypergraphs.

Corollary 6.12 also gives us a method by which we can show someone that a given hypergraph is (k, l)-partition-connected if $k \geq l$. All we have to do is to show a (k, l)-edge-connected orientation, which is a property that can be checked in polynomial time.

6.4 Directed network design problems with orientation constraints

6.4.1 Orientation constraints and submodular flows

In [50], Khanna, Naor and Shepherd introduced a new framework called "directed network design with orientation constraints". By this framework they gave a common generalization of subgraph problems such as finding a minimum cost k-rooted-connected subgraph of a digraph (that had been solved in [28]), and orientation problems like k-rooted-connected orientation of mixed graphs, discussed in [27]. The basic problem is to find a minimum cost subgraph of a digraph that satisfies a prescribed connectivity property; however, there are also *orientation constraints*: additional constraints on some designated oppositely directed pairs of edges, which require that at most one member of the pair can belong to the chosen subgraph (the term "orientation constraint" is appropriate since a constrained pair of edges can be thought of as a single undirected edge that has to be oriented or deleted, and the two possible orientations can have different costs). One of the main results in [50] stated that for the problem of finding a minimum cost subgraph that satisfies the orientation constraints and covers a given positively intersecting supermodular requirement function, the natural LP relaxation defines an integral polyhedron (note that for crossing supermodular requirement functions, this would include NP-complete problems).

In this section we extend this result to hypergraphs, and in addition show that the LP relaxation they used is in fact a TDI system; this latter result also enables us to formulate a min-max theorem. First, we show that if the requirement function is intersecting supermodular, then the orientation constraints can be incorporated into a construction of Schrijver [66] that transforms the problem without orientation constraints into a submodular flow problem. Moreover, this construction can be easily extended to the following hypergraph problem. A *mixed hypergraph* is a triple $M = (V; \mathcal{E}, \mathcal{A})$, where $\mathcal{E}$ is a set of hyperedges and $\mathcal{A}$ is a set of hyperarcs. An *oriented sub-hypergraph* of M is a sub-hypergraph of a directed hypergraph obtained from M by orienting the hyperedges in $\mathcal{E}$.

Theorem 6.13. *Let $M = (V; \mathcal{E}, \mathcal{A})$ be a mixed hypergraph, and $p : 2^V \to \mathbb{Z} \cup \{-\infty\}$ an intersecting supermodular set function. Suppose that a cost is assigned to each hyperarc in $\mathcal{A}$, and to each possible orientation of every hyperedge in $\mathcal{E}$. Then the problem of finding a minimum cost oriented sub-hypergraph of M covering p can be formulated as a submodular flow problem, solvable in polynomial time.*

Proof. The proof is a straightforward adaptation of a construction of Schrijver [66]. We define a directed bipartite graph $D = (U, V; B)$ with edge costs, where the nodes of U

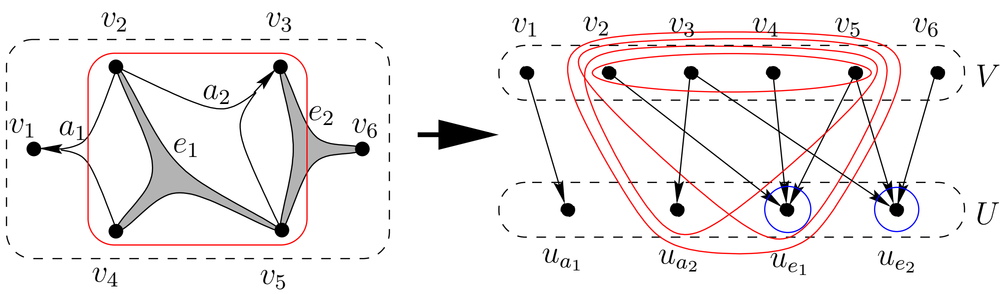

Figure 6.2: Transformation into a submodular flow problem in Theorem 6.13

correspond to the hyperedges and hyperarcs in $\mathcal{E} \cup \mathcal{A}$; we denote a node corresponding to a hyperedge or hyperarc e by u_e. The edge set B contains an edge from the head of e to u_e (with edge cost equal to the cost of e) if e is a hyperarc; if e is a hyperedge, then B contains edges from every node of e to u_e (each with cost equal to the cost of the corresponding orientation of e). A set function q is defined on the ground set $U \cup V$ as follows:

$$q(X) := \begin{cases} p(X \cap V) & \text{if } u_e \in X \text{ implies that the nodes of } e \text{ are in } X, \\ -1 & \text{if } X = \{u_e\} \text{ for some } e \in \mathcal{E}, \\ -\infty & \text{otherwise.} \end{cases}$$

Figure 6.2 shows this construction. Let $f \equiv 0$ be the lower capacity of the edges and $g \equiv 1$ the upper capacity. It is easy to see that if p is intersecting supermodular, then q is crossing supermodular. So Theorem 2.28 implies that the problem of finding a minimum cost directed subgraph $D' = (U, V; B')$ of D that satisfies

$$\delta_{D'}(X) - \varrho_{D'}(X) \geq q(X) \quad \text{for every } X \subseteq U \cup V \tag{6.6}$$

can be solved in polynomial time (note that this is a one-way submodular flow system). Since $q(\{u_e\}) = -1$ if $e \in \mathcal{E}$, u_e is the head of at most one edge of D'. Thus the subgraph D' corresponds to an oriented sub-hypergraph M' of M. It is easy to check that M' covers the requirement function p if and only if D' satisfies (6.6). □

6.4.2 TDI property

The above construction shows that when the requirement function is intersecting supermodular, the directed network design problem with orientation constraints can be transformed into a submodular flow problem which is TDI. We now prove that the linear system that we naturally associate to the original problem is also TDI, even for positively intersecting supermodular requirement functions. To formulate the appropriate linear program, the hypergraph analogue of orientation constraints must be defined. The elements of a set $\mathcal{A}'$

of hyperarcs are called *semi-parallel* if every hyperarc in $\mathcal{A}'$ is an orientation of the same hyperedge. In Theorem 6.13, we would obtain an equivalent problem if we replaced every hyperedge of the mixed graph by a set of semi-parallel hyperarcs, consisting of all possible orientations of that hyperedge, and imposed the additional constraint that at most one of these semi-parallel hyperarcs can be in the chosen sub-hypergraph. This concept of orientation constraints can be further generalized: we allow arbitrary disjoint sets of semi-parallel hyperarcs, and arbitrary lower and upper bounds on the number of hyperarcs selectable from such a set.

Theorem 6.14. *Let $D = (V, \mathcal{A})$ be a directed hypergraph, with $f : \mathcal{A} \to \mathbb{Z}_+$ and $g : \mathcal{A} \to \mathbb{Z}_+$ lower and upper integral capacities on the hyperarcs. Let $\mathcal{A}_1, \dots, \mathcal{A}_t \subseteq \mathcal{A}$ be disjoint sets of semi-parallel hyperarcs, with corresponding lower and upper bounds l_i, u_i $(i = 1, \dots, t)$. Let furthermore $p : 2^V \to \mathbb{Z}_+$ be a positively intersecting supermodular set function, and $c : \mathcal{A} \to \mathbb{Z}$ a cost function. Then the system*

$$(\mathbf{S}) \qquad \min \sum_{a \in \mathcal{A}} c(a)x(a) \tag{6.7}$$

$$\varrho_x(Z) \geq p(Z) \qquad \textit{for every } Z \subseteq V \tag{6.8}$$

$$f(a) \leq x(a) \leq g(a) \qquad \textit{for every } a \in \mathcal{A} \tag{6.9}$$

$$l_i \leq \sum_{a \in \mathcal{A}_i} x(a) \leq u_i \qquad (i = 1, \dots, t) \tag{6.10}$$

is TDI. Moreover, the values of an optimal dual solution corresponding to the inequalities (6.8) *may be assumed to be positive only on a laminar family of sets.*

Proof. Let $c : \mathcal{A} \to \mathbb{Z}$ be an integral cost function. Let y denote the dual variables associated with the inequalities in (6.8), and let z denote the dual variables associated with the other inequalities. For a hyperarc $a \in \mathcal{A}$, the dual constraints are of the form

$$\left(\sum_{X:\ a \in \Delta^-_D(X)} y(X) \right) + zb_a \leq c(a) \tag{6.11}$$

for an appropriate vector b_a. For an appropriate vector b, the dual objective function is

$$\max \left(\sum_{X \subseteq V} y(X)p(X) + zb \right). \tag{6.12}$$

Let $(y^*, z^*) \geq 0$ be an optimal dual solution such that $\sum_{Z \subseteq V} y^*(Z)$ is minimal. The main observation is that we can assume that y^* is positive only on a laminar family $\mathcal{F}$. If $p(X) = 0$, then $y^*(X) = 0$, otherwise we could decrease $y^*(X)$ to 0. Suppose that y^* is

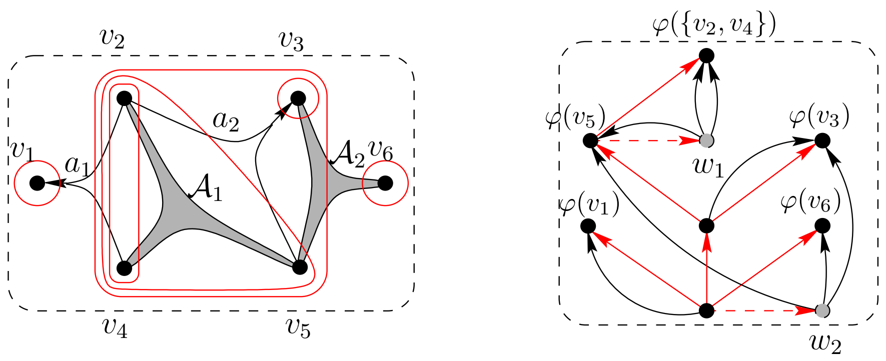

Figure 6.3: Construction of the network matrix in Theorem 6.14

positive on two intersecting sets X and Y where $p(X), p(Y) > 0$; let $\alpha = \min\{p(X), p(Y)\}$. Decrease $y^*(X)$ and $y^*(Y)$ by α, and increase $y^*(X \cap Y)$ and $y^*(X \cup Y)$ by α. Since $\varrho_a(X) + \varrho_a(Y) \geq \varrho_a(X \cap Y) + \varrho_a(X \cup Y)$ for each hyperarc a, the inequality (6.11) is preserved. The positively intersecting supermodularity of p implies that the dual objective function (6.12) does not decrease. By Claim 2.5, after a finite number of uncrossing steps, we obtain an optimal dual solution where y^* is positive on a laminar family $\mathcal{F}$.

Modify the system (**S**) by replacing (6.8) with

$$\varrho_x(Z) \geq p(Z) \text{ for every } Z \in \mathcal{F}; \tag{6.13}$$

let us denote this system by ($\mathbf{S}'$). Then (y^*, z^*) remains a feasible dual solution, and it is obviously optimal. Thus if the modified system has an integral optimal dual solution, it is optimal for the dual of (**S**) as well. The rest of the proof consists of showing that the system ($\mathbf{S}'$) can be described by a network matrix, hence it has an integral dual optimal solution by Proposition 2.25.

The construction of the corresponding network is shown on Figure 6.3. The rows of the network matrix will correspond to the edges of a directed tree $T' = (W', A'_1)$, and the corresponding lower and upper bounds will be denoted by l' and u'. The laminar family $\mathcal{F}$ has a tree-representation (T, φ) where $T = (W, A_1)$ is an arborescence; T' will include T as a subtree. For an edge $e \in A_1$ let $l'(e) = -\infty$ and $u'(e) = -p(\varphi^{-1}(W_e))$, where W_e is the component of $T - e$ entered by e. The node set W' is obtained by adding new nodes w_i $(i = 1, \ldots, t)$ to W (that is, one new node w_i for each orientation constraint set $\mathcal{A}_i$ which consists of semi-parallel hyperedges). For a set $Z \subseteq V$ let $w_Z \in W$ denote the root node of the minimal subtree of T containing all nodes of $\varphi(Z)$. To finish the construction of T', add an edge $e_i = w_{Z_i} w_i$ to A'_1 for $i = 1, \ldots, t$, where Z_i is the node set of the hyperedge whose orientations are in $\mathcal{A}_i$. Define the corresponding lower and upper bounds as $l'(e_i) = l_i$, $u'(e_i) = u_i$.

The columns of the matrix will represent a set A_2' of directed edges, with a one-to-one correspondence between the hyperarcs in $\mathcal{A}$ and the edges in A_2'. To a hyperarc $a \in \mathcal{A}_i$, assign an edge $w_i h(a)$. To hyperarcs $a \in \mathcal{A}$ not appearing in any $\mathcal{A}_i$, assign an edge $w_Z h(a)$, where Z is the node set of a.

Let N denote the network matrix given by the above network $(W'; A_1', A_2')$. Then by Proposition 2.25 the system

$$\{x : A_2' \to \mathbb{Q} :\ l' \leq Nx \leq u',\ f \leq x \leq g\}$$

is TDI. Moreover, by the one-to-one correspondence between the edges in A_2' and the hyperarcs in $\mathcal{A}$, this system is equivalent to the system $(\mathbf{S}')$. This implies that $(\mathbf{S}')$ has an integral dual optimal solution, which in turn is an optimal dual solution for $(\mathbf{S})$. □

6.4.3 Min-max theorems

Theorem 6.14 implies that the polyhedron described by $(\mathbf{S})$ is integral, and for every integer cost function there exists an integral optimal dual solution where the family of the sets with positive dual variable is laminar. This allows us to formulate fairly friendly new min-max formulas for some graph problems. For example, what is the maximum number of undirected edges, or the maximum number of directed edges, that can be removed from a mixed graph such that the obtained subgraph has an orientation covering a given intersecting supermodular set function p? The following corollary describes a min-max formula that involves both of these problems.

Corollary 6.15. *Let $G = (V; E, A)$ be a mixed graph (where E is the set of undirected edges and A is the set of directed edges). Let $c : E \cup A \to \{0, 1\}$ be a cost function, and $p : 2^V \to \mathbb{Z}_+$ a positively intersecting supermodular set function. Then the minimum cost of a subgraph that has an orientation covering p equals*

$$\max_{\mathcal{F} \text{ laminar}} \left(\sum_{X \in \mathcal{F}} p(X) - e_E(\mathcal{F}) - \sum_{X \in \mathcal{F}} \varrho_A(X) + \mu(\mathcal{F}) \right), \tag{6.14}$$

where $\mu(\mathcal{F})$ is the sum of the costs of the undirected and directed edges that enter at least one member of $\mathcal{F}$.

Proof. To formulate the problem in the terms of Theorem 6.14, let the directed hypergraph $D = (V, \mathcal{A})$ be the digraph obtained from G by replacing the undirected edges of G by a pair of oppositely directed edges, and let us assign an orientation constraint to every such pair, with bounds $l_i := 0$ and $u_i := 1$. Let the cost of the edges in a pair be the cost of the corresponding undirected edge in E. The capacities of the edges are bounded by $f \equiv 0$, $g \equiv 1$.

For a $\{0,1\}$-valued cost function c, consider the system ($\mathbf{S}$), and let the dual solutions be denoted by (y, z), where y consists of the dual variables associated to the constraints in (6.8). Take an integral dual optimal solution $(y^*, z^*) \geq 0$, where y^* is positive on a laminar family, and $|z^*|$ is minimal. Let $\mathcal{F}$ be the laminar family where every set X has multiplicity $y^*(X)$. It follows from the minimality of $|z^*|$ that the objective value of this dual solution is

$$\sum_{X\in\mathcal{F}} p(X) - \sum_{e\in E} (e_e(\mathcal{F}) - c(e))^+ - \sum_{a\in A} \left(\sum_{X\in\mathcal{F}} \rho_a(X) - c(a) \right)^+ = \\ = \sum_{X\in\mathcal{F}} p(X) - e_E(\mathcal{F}) - \sum_{X\in\mathcal{F}} \varrho_A(X) + \mu(\mathcal{F}).$$

Conversely, the value of (6.14) corresponds to the value of the following dual solution (y^*, z^*). Let $\mathcal{F}$ be a laminar family where the maximum is attained in (6.14). For $X \subseteq V$, let $y^*(X)$ be the multiplicity of X in $\mathcal{F}$. Define the values of the dual variables in z^* as required by the dual constraints, always setting a variable corresponding to an edge to 0 if the edge also belongs to an orientation constraint. In this case the dual objective value is equal to the expression (6.14). □

6.5 Local connectivity requirements

Theorems 6.13 and 6.14 show that in relation to the covering of intersecting supermodular set functions, a relatively large class of orientation problems can be efficiently solved. The preceding results of the chapter indicate that a somewhat more restricted class is solvable concerning the covering of crossing supermodular set functions. In this section we attempt to further relax the condition of crossing supermodularity. As we have mentioned at the beginning of the chapter, it is open whether some analogue of Theorem 6.4 can be proved for hypergraphs. The problems discussed in this section are much less ambitious, but nevertheless they have some interesting corollaries.

6.5.1 Local requirement for one pair of nodes

We consider k-edge-connected orientations of graphs and hypergraphs, where the number of edge-disjoint paths required between two designated special nodes may be more than k. First we formulate a partition-type condition for the hypergraph case, and prove its sufficiency using a modified uncrossing method.

Theorem 6.16. *Let $H = (V, \mathcal{E})$ be a hypergraph, $s, t \in V$ designated nodes, and $k_1, k_2 \geq k$ positive integers. H has a k-edge-connected orientation such that there are k_1 edge-disjoint paths from s to t and k_2 edge-disjoint paths from t to s if and only if*

$$e_H(\mathcal{F}) \geq \sum_{X \in \mathcal{F}} p(X) \tag{6.15}$$

for every partition $\mathcal{F}$, where

$$p(X) := \begin{cases} 0 & \text{if } X = \emptyset \text{ or } X = V, \\ k_1 & \text{if } X \text{ is an } \overline{s}t\text{-set}, \\ k_2 & \text{if } X \text{ is a } \overline{t}s\text{-set}, \\ k & \text{otherwise.} \end{cases} \tag{6.16}$$

Proof. We may assume that every hyperedge of H contains at least 2 nodes. The goal is to find an orientation of H that covers p. Observe that the set function p has none of the properties discussed in the previous sections (it is not crossing supermodular, monotone decreasing, or symmetric). As in the proof of Theorem 6.11, we increase the value of p on the singletons so that every singleton $\{v\}$ is in a tight partition $\mathcal{F}_v$ (a partition that satisfies (6.15) by equality); let $\mathcal{F} := \sum_{v \in V} \mathcal{F}_v$ be the sum of these partitions, and let p' denote the modified set function; then

$$\sum_{X \in \mathcal{F}} p'(X) = e_H(\mathcal{F}). \tag{6.17}$$

Apply one of the following three operations on $\mathcal{F}$ as long as any of them can be applied (see Figure 6.4):

(i) Uncross $X \in \mathcal{F}$ and $Y \in \mathcal{F}$ (see Lemma 2.4) if they are crossing unless one of them is an $\overline{s}t$-set and the other is a $\overline{t}s$-set;

(ii) If $\mathcal{F}$ contains a co-partition, replace it by the partition obtained by taking the complement of every member;

(iii) If $X \in \mathcal{F}$ is an $\overline{s}t$-set, $Y \in \mathcal{F}$ is a $\overline{t}s$-set, and there is a sub-family $\mathcal{G} \subseteq \mathcal{F}$ such that $\mathrm{co}(\mathcal{G})$ is a partition of $X \cap Y$, replace X, Y and $\mathcal{G}$ in $\mathcal{F}$ by $X - Y$, $Y - X$ and $\mathrm{co}(\mathcal{G})$.

Claim 6.17. *These operations do not increase $e_H(\mathcal{F})$, and do not decrease $\sum_{X \in \mathcal{F}} p'(X)$.*

Proof. A simple case analysis shows that the operations do not increase $e_H(\mathcal{F})$, as it suffices to check that the operations do not increase $\sum_{X \in \mathcal{F}} \varrho_a(X)$ for any hyperarc a. An even more simple case analysis shows that the operations do not decrease $\sum_{X \in \mathcal{F}} p(X)$, consequently they cannot decrease the value $\sum_{X \in \mathcal{F}} p'(X)$, since singletons are never removed from the family. □

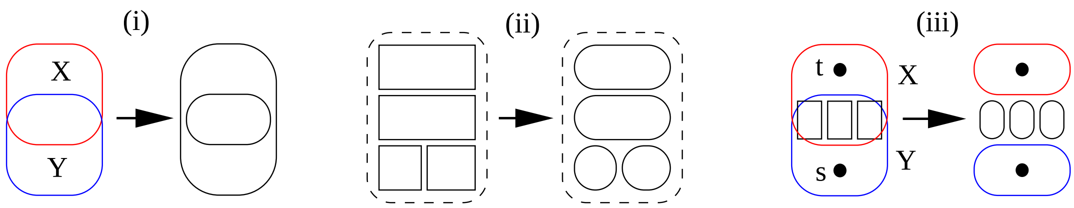

Figure 6.4: The 3 operations in the proof of Theorem 6.16. The non-rounded rectangles represent their complement.

Obviously $\mathcal{F}$ remains regular throughout the process. The second and the third operations decrease $h_\emptyset(\mathcal{F})$, and by Claim 2.5 the first operation can be applied only finitely many times consecutively, so after a finite number of steps none of the three operations can be applied. Let us denote the obtained regular family by $\mathcal{F}'$; (6.17) and Claim 6.17 imply that $\sum_{X\in\mathcal{F}'} p'(X) \geq e_H(\mathcal{F}')$. Let $\mathcal{F}''$ be the regular family obtained from $\mathcal{F}'$ by decreasing the multiplicity of every singleton by one.

Claim 6.18. *$\mathcal{F}''$ decomposes into partitions.*

Proof. We can assume that there is an $\overline{s}t$-set and a $\overline{t}s$-set, otherwise the unavailability of the first and the second operation would imply that $\mathcal{F}''$ is a cross-free family that decomposes into partitions. The $\overline{s}t$-sets in $\mathcal{F}''$ form a chain, the $\overline{t}s$ sets likewise. Let X be the minimal $\overline{s}t$-set, and Y the maximal $\overline{t}s$-set in $\mathcal{F}''$.

If $X \cap Y \neq \emptyset$, then for every $v \in X \cap Y$ there is a $\overline{v}t$-set in $\mathcal{F}''$, since $\mathcal{F}''$ is regular; let $\mathcal{B}$ denote the family of these sets. By the minimality of X, the members of $\mathcal{B}$ are not $\overline{s}t$-sets. Furthermore, they are neither crossing each other, nor X, nor Y, since the first operation cannot be applied. This is only possible if the minimal sets in $\text{co}(\mathcal{B})$ define a partition of $X\cap Y$. But then the third operation would have been applicable, contradicting the assumption.

Thus X and Y are disjoint. For every $v \in V - X - Y$ there is a $\overline{t}v$-set in $\mathcal{F}''$, since $\mathcal{F}''$ is regular. By the maximality of Y, these sets are not $\overline{t}s$-sets, so they are disjoint from X and Y, otherwise they would cross X or Y. The minimal such sets are also disjoint from each other, so they form a partition $\mathcal{P}$ of $V - X - Y$. Thus $\mathcal{F}''$ contains the partition $\mathcal{P} + \{X, Y\}$. By induction, $\mathcal{F}''$ decomposes into partitions. $\square$

If the conditions of the theorem are met, then $\sum_{X\in\mathcal{F}''} p'(X) \leq e_H(\mathcal{F}'')$ by Claim 6.18, so $\sum_{X\in\mathcal{F}'} p'(X) \geq e_H(\mathcal{F}')$ implies that the partition formed by the singletons must be tight. Let the vector $x : V \to \mathbb{Z}$ be defined by $x(v) := p'(\{v\})$; then $x(V) = |\mathcal{E}|$.

The end of the proof is the same as for Theorem 6.11. We define the partition $\mathcal{F}_Y := \{Y\} \cup \{\{v\} :\ v \in V - Y\}$ for every set $Y \subset V$. The conditions of the theorem imply that

$$x(Y) = |\mathcal{E}| - x(V - Y) = |\mathcal{E}| + p'(Y) - \sum_{Z \in \mathcal{F}_Y} p'(Z)$$
$$\geq |\mathcal{E}| - e_H(\mathcal{F}_Y) + p'(Y) = i_H(Y) + p'(Y).$$

Thus $x(Y) \geq i_H(Y) + p'(Y)$ for every set $Y \subseteq V$, and by Lemma 6.10 there is an orientation with in-degree vector x that covers p', hence it covers p. □

6.5.2 Characterization of (k, l)-partition-connectivity $(k < l)$

An interesting corollary of Theorem 6.16 is a characterization of (k, l)-partition-connected hypergraphs for $k < l$.

Corollary 6.19. *A hypergraph $H = (V, \mathcal{E})$ is (k, l)-partition-connected for $k < l$ if and only if for every pair $s, t \in V$ it has a k-edge-connected orientation with l edge-disjoint paths from s to t.*

Proof. It is easy to check that if for every pair $s, t \in V$ the hypergraph H has a k-edge-connected orientation with l edge-disjoint paths from s to t, then $e_H(\mathcal{F}) \geq k(|\mathcal{F}| - 1) + l$ must hold for every nontrivial partition $\mathcal{F}$, since we can always choose s and t to be in different members of the partition.

To prove the other direction, let H be a (k, l)-partition-connected hypergraph. Then Theorem 6.16 implies that for every pair $s, t \in V$ it has a k-edge-connected orientation with l edge-disjoint paths from s to t. □

Corollary 6.19 gives a method (albeit a relatively complicated one) for showing someone that a given hypergraph is (k, l) partition-connected for some $k < l$. For every pair of nodes $s, t \in V$, we can give a k-edge-connected orientation that contains l edge-disjoint paths from s to t, and this property can be verified in polynomial time.

6.5.3 Characterization of $(2k + 1)$-edge-connected graphs

A simple observation shows that for graphs the condition of Theorem 6.16 can be further simplified: instead of partitions, it suffices to consider cut conditions.

Theorem 6.20. *Let $G = (V, E)$ be an undirected graph, with $s, t \in V$ designated nodes, and let k, k_1, k_2 be positive integers for which $k_1, k_2 \geq k$. Then G has a k-edge-connected orientation such that there are k_1 edge-disjoint paths from s to t and k_2 edge-disjoint paths* *from t to s if and only if $d_G(X) \geq 2k$ for every $\emptyset \neq X \subset V$, and $d_G(X) \geq k_1 + k_2$ for every $\overline{s}t$-set.*

Proof. Let us define p by (6.16). Suppose indirectly that the conditions of the theorem hold, yet there is a partition $\mathcal{F}$ such that

$$e_G(\mathcal{F}) = \sum_{X \in \mathcal{F}} \frac{d_G(X)}{2} < \sum_{X \in \mathcal{F}} p(X).$$

This implies that there is a member X of $\mathcal{F}$ such that $p(X) > d_G(X)/2$, and X must separate s and t, otherwise it would violate the conditions. Let Y be the other member of $\mathcal{F}$ separating s and t. Then either $k_1 + k_2 > (d_G(X) + d_G(Y))/2$, or $k > d_G(Z)/2$ for some other member Z of $\mathcal{F}$, contradicting the conditions. □

Theorem 6.20 can also be proved using a different approach that does not seem to extend to hypergraphs, namely a simple application of Theorem 3.3 of Mader on undirected splitting-off preserving local edge-connectivity.

Alternative proof of Theorem 6.20. We use induction on the number of edges of G. Call a set X *tight* if $d_G(X) = k_1 + k_2$ and X separates s and t, or $d_G(X) = 2k$ and X does not separate s and t. We may assume that every edge enters a tight set. If every edge of G enters a tight set separating s and t, then the edge set of G can be partitioned into $k_1 + k_2$ simple paths between s and t, and every node v is reached by at least k such paths, since $d_G(v) \geq 2k$. Let $\vec{G}$ be the digraph obtained by orienting k_1 paths from s to t, and k_2 paths from t to s; then $\varrho_{\vec{G}}(X) \geq k$ and $\varrho_{\vec{G}}(V - X) \geq k$ for every set $\emptyset \neq X \subseteq V - \{s,t\}$, hence $\vec{G}$ is a good orientation.

We can now assume that there exists a non-empty minimal tight set W not containing s and t. Observe that if X and Y are crossing tight sets, then either one of them is an $\overline{s}t$-set and the other is a $\overline{t}s$-set, or $X \cap Y$ is tight; therefore the minimality of W implies that $i_G(W) = 0$, since an edge spanned by W would not enter a tight set. Thus W is a singleton $\{w\}$ with $d_G(w) = 2k$.

By Theorem 3.3, there is a complete splitting at w that preserves the conditions of the Theorem. By induction, the resulting graph has a good orientation, which can be transformed into a good orientation of the original graph by the inverse operation of edge splitting: if the orientation of the edge that resulted from the splitting-off is a directed edge uv, then replace it by directed edges uw and wv. □

Finally, let us mention a corollary regarding $(2k+1)$-edge-connected graphs. This can be seen as a counterpart of Theorem 6.1 of Nash-Williams on the characterization of $2k$-edge-connected graphs, although it is much less elegant. Theorem 6.20 implies the following:

Corollary 6.21. *A graph $G = (V, E)$ is $(2k+1)$-edge-connected if and only if for every pair $s, t \in V$ it has a k-edge-connected orientation with $k+1$ edge-disjoint paths from s to t.* □

Chapter 7

Combined augmentation and orientation

7.1 Introduction

The problems discussed in this chapter have ties to both connectivity orientation and connectivity augmentation. The results of Chapter 6 suggest that the property that a hypergraph has an orientation with high edge-connectivity may itself be considered as a connectivity property. For example, the existence of a (k,l)-edge-connected orientation is equivalent to the (k,l)-partition-connectivity of the hypergraph (defined in Chapter 5). It is reasonable to ask when can connectivity augmentation be solved for this kind of connectivity properties. In other words, we examine the solvability of the following problem: given a hypergraph H_0, add an optimal set of hyperedges to H_0 such that the resulting hypergraph has an orientation with some prescribed connectivity property.

At first glance this question could seem to be closely related to the undirected augmentation problems described in Chapter 3. Interestingly, it turns out that it has more in common with the connectivity augmentation of directed hypergraphs, discussed in Chapter 4. In particular, we will extensively use the directed splitting-off technique described there.

Problems will be presented in the usual framework involving the covering of set functions, but in the first sections we restrict ourselves to non-negative crossing supermodular set functions. This class includes the requirement function for (k,l)-edge-connected orientations, and the obtained results give min-max formulas on (k,l)-partition-connectivity augmentation.

The final section contains some results on combined augmentation-orientation problems when the requirement function is only positively crossing supermodular. The characterizations become much more complicated in that case, and the solvability of minimum cardinality augmentation remains open; we only address the degree-specified problem.

The new results presented here are partly from [35], a joint paper with András Frank, and partly from [52], a joint work with Márton Makai. It should be noted that the theorems of the chapter give new results even in the special case when they are restricted to graphs. For example, a min-max theorem on minimum cardinality (k, l)-partition-connectivity augmentation of graphs can be derived from Corollary 7.11.

7.2 Augmentation to meet orientability requirements

7.2.1 Degree specified augmentation

As in the case of undirected and directed edge-connectivity augmentation, we first prove a degree specified augmentation result. The conditions on the requirement function are a bit more restrictive now, since we require non-negativity, like in Theorem 6.11. Note that Theorem 6.11 actually corresponds to the special case of the following theorem when the degree specification is 0 on every node.

We have seen that there can be different objectives in an augmentation problem, depending on the restrictions on the sizes of the added hyperedges. The results presented here give characterizations for both the uniform and the unrestricted augmentation problem. The given characterizations are "good" in the sense that they provide an easily verifiable certificate if the augmentation is impossible.

Theorem 7.1. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $p : 2^V \to \mathbb{Z}_+$ a non-negative crossing supermodular set function, $m : V \to \mathbb{Z}_+$ a degree specification, and $0 \leq \gamma \leq m(V)/2$ an integer. There exists a hypergraph $H = (V, \mathcal{E})$ with γ hyperedges such that $H_0 + H$ has an orientation covering p and $d_H(v) = m(v)$ for every $v \in V$ if and only if the following hold for every partition $\mathcal{F}$ of V:*

$$\gamma \geq \sum_{Z \in \mathcal{F}} p(Z) - e_{H_0}(\mathcal{F}), \tag{7.1}$$

$$\min_{X \in \mathcal{F}} m(V - X) \geq \sum_{Z \in \mathcal{F}} p(Z) - e_{H_0}(\mathcal{F}), \tag{7.2}$$

$$\min_{\mathcal{F}' \subseteq \mathcal{F},\ X = \cup \mathcal{F}'} (m(V - X) + (|\mathcal{F}'| - 1)\gamma) \geq \sum_{Z \in \mathcal{F}} p(V - Z) - e_{H_0}(\mathrm{co}(\mathcal{F})). \tag{7.3}$$

In addition, H can be chosen so that

$$\left\lfloor \frac{m(V)}{\gamma} \right\rfloor \leq |e| \leq \left\lceil \frac{m(V)}{\gamma} \right\rceil \quad \textit{for every } e \in \mathcal{E}. \tag{7.4}$$

Proof. The right hand side of the inequalities is the deficiency of the hyperedges of H_0 on the appropriate families. The necessity of the conditions follows from the observation that the left hand side is always an upper bound on the contribution of the new hyperarcs. In (7.1): every new hyperarc can enter at most one set of $\mathcal{F}$; in (7.2): every hyperarc that enters a set of $\mathcal{F}$ must have a node in $V - X$; in (7.3): the number of sets of $\mathrm{co}(\mathcal{F})$ that a new hyperarc enters is at most $|\mathcal{F}'| - 1$ plus the number of nodes it has in $V - X$.

The proof of sufficiency is based on the orientation of an extended hypergraph. We add a new node z to the set of nodes, and for every $v \in V$ we add $m(v)$ parallel edges between v and z; the resulting hypergraph is denoted by $H_0' = (V', \mathcal{E}_0')$. Our first aim is to find an orientation $\vec{H}_0'$ of H_0' that has the following properties:

$$\varrho_{\vec{H}_0'}(V) = \gamma, \tag{7.5}$$

$$\varrho_{\vec{H}_0'}(X) \geq p(X) \quad \text{if } \emptyset \neq X \subset V, \tag{7.6}$$

$$\varrho_{\vec{H}_0'}(X + z) \geq p(X) \quad \text{if } \emptyset \neq X \subset V. \tag{7.7}$$

To find such an orientation, we use Lemma 6.10. A vector $x : V \to \mathbb{Z}_+$ is called *feasible* if it is the vector of in-degrees (restricted to V) of an orientation satisfying (7.5)–(7.7). It is easy to see using Lemma 6.10 that x is feasible if and only if $x(V) = |\mathcal{E}_0| + \gamma$ and $x(Z) \geq p_m(Z)$ for every $Z \subseteq V$, where

$$p_m(X) := p(X) + i_{H_0}(X) + (\gamma - m(V - X))^+ \qquad (X \subseteq V). \tag{7.8}$$

Thus a vector x is feasible if and only if it is an integral element of $B(p_m)$ (as defined in (2.5)).

Claim 7.2. *The set function p_m is crossing supermodular.*

Proof. The crossing supermodularity of p and (1.3) implies that $p + i_{H_0}$ is crossing supermodular. Let $m^*(X) := (\gamma - m(V - X))^+$; we show that this set function is fully supermodular. If $m^*(Y) = 0$, then $m^*(X) + m^*(Y) = m^*(X) \leq m^*(X \cup Y) = m^*(X \cap Y) + m^*(X \cup Y)$. If $m^*(X), m^*(Y) > 0$, then $m^*(X) + m^*(Y) = 2\gamma - 2m(V) + m(X \cap Y) + m(X \cup Y) \leq m^*(X \cap Y) + m^*(X \cup Y)$. The sum of a crossing supermodular and a fully supermodular function is crossing supermodular. □

Theorem 2.11 implies that if $B(p_m)$ is non-empty, then it is a base polyhedron, hence its vertices are integral.

Claim 7.3. *If conditions (7.1)-(7.3) are satisfied, then $B(p_m)$ is non-empty.*

Proof. By Theorem 2.11, it suffices to show that

$$\sum_{X\in\mathcal{F}} p_m(X) \leq |\mathcal{E}_0| + \gamma, \tag{7.9}$$

$$\sum_{X\in\mathcal{F}} p_m(V-X) \leq (|\mathcal{F}|-1)(|\mathcal{E}_0| + \gamma) \tag{7.10}$$

both hold for every partition $\mathcal{F}$. Note that $\gamma - m(V-X)$ can be positive for at most one member of a partition, since $\gamma \leq m(V)/2$. Thus (7.9) follows from (1.12), and either (7.2) or (7.1), depending on whether $\mathcal{F}$ has such a member or not. The inequality (7.10) follows from (1.12) and (7.3), if $\mathcal{F}'$ is chosen to consist of the members X of $\mathcal{F}$ for which $\gamma - m(X) \leq 0$. □

By Theorem 2.11, $B(p_m)$ is a base polyhedron with integral vertices, and any such vertex x is the vector of in-degrees (restricted to V) of an orientation $\vec{H}_0'$ of H_0' satisfying (7.5)–(7.7).

Let $m_i(v)$ be the multiplicity of the edge zv in $\vec{H}_0'$, $m_o(v)$ be the multiplicity of the edge vz in $\vec{H}_0'$, and let $\vec{H}_0$ denote the directed hypergraph obtained from $\vec{H}_0'$ by deleting the node z. Then $m_i(X) \geq p(X) - \varrho_{\vec{H}_0}(X)$ and $m_o(V-X) \geq p(X) - \varrho_{\vec{H}_0}(X)$ for every $X \subseteq V$. By (1.4) and the crossing supermodularity of p, the set function $q(X) := p(X) - \varrho_{\vec{H}_0}(X)$ is crossing supermodular. Theorem 4.7 asserts the existence of a directed hypergraph $D = (V, \mathcal{A})$ that covers q, and satisfies the degree specifications m_i and m_o. This means that $\vec{H}_0 + D$ covers p, and the undirected hypergraph H that underlies D satisfies the degree specification m. Theorem 4.7 also ensures that (7.4) is satisfied. Since $\vec{H}_0 + D$ is an orientation of $H_0 + H$, this completes the proof of Theorem 7.1. □

If the requirement function p is monotone decreasing or symmetric, then the conditions of Theorem 7.1 can be simplified.

Theorem 7.4. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $p : 2^V \to \mathbb{Z}_+$ a monotone decreasing or symmetric non-negative crossing supermodular set function, $m : V \to \mathbb{Z}_+$ a degree specification and $0 \leq \gamma \leq m(V)/2$ an integer. There exists a hypergraph H with γ hyperedges satisfying the degree-specification m such that $H_0 + H$ has an orientation covering p if and only if the following hold for every partition $\mathcal{F}$ of V:*

$$\gamma \geq \sum_{Z\in\mathcal{F}} p(Z) - e_{H_0}(\mathcal{F}), \tag{7.11}$$

$$\min_{X\in\mathcal{F}} m(V-X) \geq \sum_{Z\in\mathcal{F}} p(Z) - e_{H_0}(\mathcal{F}). \tag{7.12}$$

In addition, H can be chosen so that

$$\left\lfloor \frac{m(V)}{\gamma} \right\rfloor \leq |e| \leq \left\lceil \frac{m(V)}{\gamma} \right\rceil \quad \text{for every } e \in \mathcal{E}. \tag{7.13}$$

Proof. By definition, $e_{H_0}(\mathcal{F}) \leq e_{H_0}(\text{co}(\mathcal{F}))$ for every partition $\mathcal{F}$ of V, and the monotonicity or symmetry of p implies that $\sum_{Z\in\text{co}(\mathcal{F})} p(Z) \leq \sum_{Z\in\mathcal{F}} p(Z)$ also holds. It is easy to see from this that (7.3) is implied by (7.1) if $|\mathcal{F}'| = 0$ or $|\mathcal{F}'| \geq 2$, and it is implied by (7.2) if $|\mathcal{F}'| = 1$. □

7.2.2 Minimum cardinality augmentation

As in the case of augmentation problems considered in Chapters 3 and 4, the characterization of the degree specifications that allow a good augmentation helps to deduce a characterization of the minimum number of hyperedges needed. In the present case we obtain the following theorem:

Theorem 7.5. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $p : 2^V \to \mathbb{Z}_+$ a non-negative crossing supermodular set function, $\sigma \geq 0$ and $0 \leq \gamma \leq \sigma/2$ integers. There exists a hypergraph $H = (V, \mathcal{E})$ with γ hyperedges of total size σ such that $H_0 + H$ has an orientation covering p if and only if*

$$\sigma + \gamma h_X(\mathcal{F}_1) + (\sigma - \gamma) h_X(\mathcal{F}_2) \geq \sum_{Z\in\mathcal{F}_1+\text{co}(\mathcal{F}_2)} p(Z) - e_{H_0}(\mathcal{F}_1 + \text{co}(\mathcal{F}_2)) \tag{7.14}$$

whenever $\mathcal{F}_1$ and $\mathcal{F}_2$ are tree-compositions of some $X \subseteq V$, $\mathcal{F}_1 + \mathcal{F}_2$ is cross-free, and $h_X(\mathcal{F}_2) \leq 0$ (i.e. either $\mathcal{F}_2$ is a partition of X, or $X = V$ and $\mathcal{F}_2 = \emptyset$).

In addition, H can be chosen so that

$$\left\lfloor \frac{\sigma}{\gamma} \right\rfloor \leq |e| \leq \left\lceil \frac{\sigma}{\gamma} \right\rceil \quad \text{for every } e \in \mathcal{E}. \tag{7.15}$$

Proof. The right hand side of (7.14) is the deficiency of H_0 on the family $\mathcal{F}_1 + \text{co}(\mathcal{F}_2)$. The number of sets of $\mathcal{F}_1$ that a new hyperarc enters is at most $h_X(\mathcal{F}_1)$, plus 1 if its head is in X. The number of sets of $\text{co}(\mathcal{F}_2)$ that a new hyperarc a enters is at most $(|a|-1)h_X(\mathcal{F}_2)$ plus the number of tail nodes it has in X. This shows the necessity of (7.14). To prove sufficiency, we define the following functions for every $X \subseteq V$ and compositions $\mathcal{F}_1, \mathcal{F}_2$ of X:

$$Q_X(\mathcal{F}_1, \mathcal{F}_2) := \sum_{Z\in\mathcal{F}_1+\text{co}(\mathcal{F}_2)} p(Z) - e_{H_0}(\mathcal{F}_1 + \text{co}(\mathcal{F}_2)) - \gamma h_X(\mathcal{F}_1) - (\sigma - \gamma) h_X(\mathcal{F}_2)),$$

$$\begin{aligned} q(X) := \max\{Q_X(\mathcal{F}_1, \mathcal{F}_2) :\ & \mathcal{F}_1 \text{ and } \mathcal{F}_2 \text{ are tree-compositions of } X, \\ & \mathcal{F}_1 + \mathcal{F}_2 \text{ is cross-free, } h_X(\mathcal{F}_2) \leq 0\}. \end{aligned}$$

Condition (7.14) is equivalent to the inequality $\max_{X \subseteq V} q(X) \leq \sigma$; let us assume that this holds. We may observe that if a degree-specification $m : V \to \mathbb{Z}_+$ satisfies $m(Y) \geq q(Y)$ for every $Y \subseteq V$ and $m(V) = \sigma$, then m satisfies (7.1)–(7.3). Indeed, (7.1) for a partition $\mathcal{F}$ follows if we consider $m(Y) \geq q(Y)$ with the choice of $Y = V$, $\mathcal{F}_1 = \mathcal{F}$, and $\mathcal{F}_2 = \emptyset$ (hence $h_Y(\mathcal{F}_1) = 0$ and $h_Y(\mathcal{F}_2) = -1$). Inequality (7.2) for a partition $\mathcal{F}$ and a member $X \in \mathcal{F}$ is obtained by setting $Y = V - X$, $\mathcal{F}_1 = \mathcal{F} - \{X\}$ (a partition of Y) and $\mathcal{F}_2 = \{Y\}$ (thus $h_Y(\mathcal{F}_1) = 0$ and $h_Y(\mathcal{F}_2) = 0$). Finally, (7.3) for a partition $\mathcal{F}$, a subpartition $\mathcal{F}' \subseteq \mathcal{F}$ and $X = \cup\mathcal{F}'$ is obtained from $m(Y) \geq q(Y)$ with the settings $Y = V - X$, $\mathcal{F}_1 = \text{co}(\mathcal{F}')$, and $\mathcal{F}_2 = \mathcal{F} - \mathcal{F}'$ (in which case $h_Y(\mathcal{F}_1) = |\mathcal{F}_1| - 1$ and $h_Y(\mathcal{F}_2) = 0$).

By Theorem 7.1 the existence of such a degree-specification m implies the existence of a hypergraph H that satisfies the requirements. To prove the existence of a good degree-specification, we use the properties of a set function slightly different from q:

$$q'(X) := \max\{Q_X(\mathcal{F}_1, \mathcal{F}_2) : \ \mathcal{F}_1 \text{ and } \mathcal{F}_2 \text{ are tree-compositions of } X\}.$$

Claim 7.6. *The value $Q_X(\mathcal{F}_1, \mathcal{F}_2)$ does not decrease if we remove a partition or a co-partition of V from $\mathcal{F}_1$ or $\mathcal{F}_2$.*

Proof. It is easy to see that if $X \cap Y = \emptyset$, $\mathcal{F}_1^X, \mathcal{F}_2^X$ are compositions of X, and $\mathcal{F}_1^Y, \mathcal{F}_2^Y$ are compositions of Y, then

$$Q_X(\mathcal{F}_1^X, \mathcal{F}_2^X) + Q_Y(\mathcal{F}_1^Y, \mathcal{F}_2^Y) = Q_{X \cup Y}(\mathcal{F}_1^X + \mathcal{F}_1^Y, \mathcal{F}_2^X + \mathcal{F}_2^Y). \tag{7.16}$$

The case $Y = \emptyset$ proves the claim, since $q(V) \leq \sigma$ implies that $q(\emptyset) \leq 0$. □

Claim 7.7. *The set function q' is fully supermodular.*

Proof. Let $X, Y \subseteq V$, and suppose that the maximum in the definition of q' is reached on families $\mathcal{F}_1^X, \mathcal{F}_2^X$, and $\mathcal{F}_1^Y, \mathcal{F}_2^Y$, respectively. Let $\mathcal{F}_1 := \mathcal{F}_1^X + \mathcal{F}_1^Y$, $\mathcal{F}_2 := \mathcal{F}_2^X + \mathcal{F}_2^Y$. We apply the following operations, as long as any of them is possible:

(i) If $Z_1, Z_2 \in \mathcal{F}_1$ are crossing, then replace them in $\mathcal{F}_1$ by $Z_1 \cap Z_2$ and $Z_1 \cup Z_2$.

(ii) If $Z_1, Z_2 \in \mathcal{F}_2$ are crossing, then replace them in $\mathcal{F}_2$ by $Z_1 \cap Z_2$ and $Z_1 \cup Z_2$.

Lemma 2.4 implies that after a finite number of steps, the resulting families $\mathcal{F}_1'$ and $\mathcal{F}_2'$ become cross-free. By Lemma 2.3 $\mathcal{F}_i'$ decomposes into a composition $\mathcal{F}_i^{X \cap Y}$ of $X \cap Y$ and a composition $\mathcal{F}_i^{X \cup Y}$ of $X \cup Y$ ($i = 1, 2$); and all of these families are cross-free. The crossing supermodularity of p implies that $\sum_{Z \in \mathcal{F}_1} p(Z) \leq \sum_{Z \in \mathcal{F}_1'} p(Z)$ and

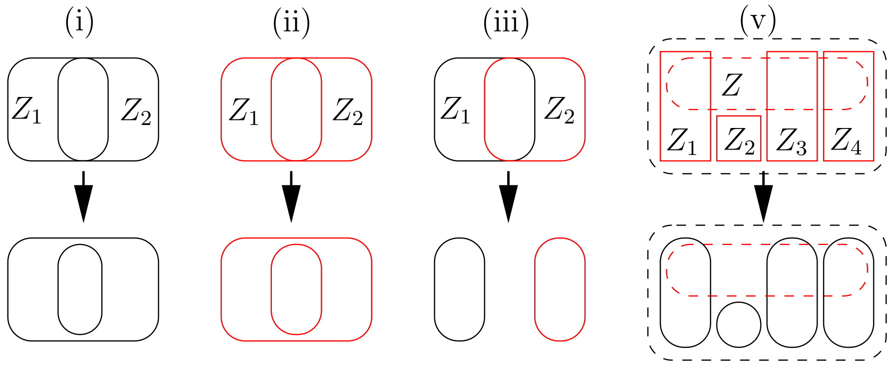

Figure 7.1: Operations (i), (ii), (iii), and (v) in the proof of Claim 7.8

$\sum_{Z\in\mathcal{F}_2} p(V-Z) \le \sum_{Z\in\mathcal{F}_2'} p(V-Z)$. It is easy to check using Proposition 1.17 that $e_{H_0}(\mathcal{F}_1^X + \mathrm{co}(\mathcal{F}_2^X)) + e_{H_0}(\mathcal{F}_1^Y + \mathrm{co}(\mathcal{F}_2^Y)) \ge e_{H_0}(\mathcal{F}_1^{X\cap Y} + \mathrm{co}(\mathcal{F}_2^{X\cap Y})) + e_{H_0}(\mathcal{F}_1^{X\cup Y} + \mathrm{co}(\mathcal{F}_2^{X\cup Y}))$. Since the uncrossing does not change the height of families, we also have $h_X(\mathcal{F}_i^X) + h_Y(\mathcal{F}_i^Y) = h_{X\cap Y}(\mathcal{F}_i^{X\cap Y}) + h_{X\cup Y}(\mathcal{F}_i^{X\cup Y})$ $(i = 1,2)$. Thus $Q(\mathcal{F}_1^X, \mathcal{F}_2^X) + Q(\mathcal{F}_1^Y, \mathcal{F}_2^Y) \le Q(\mathcal{F}_1^{X\cap Y}, \mathcal{F}_2^{X\cap Y}) + Q(\mathcal{F}_1^{X\cup Y}, \mathcal{F}_2^{X\cup Y})$. Using Claim 7.6, we obtain that $q'(X) + q'(Y) \le q'(X\cap Y) + q'(X\cup Y)$. □

Claim 7.7 and Theorem 2.11 imply that there exists a vector $m : V \to \mathbb{Z}_+$ with $m(V) = \sigma$ that satisfies $m(X) \ge q'(X)$ for every $X \subseteq V$ if and only if $\max_{X\subseteq V} q'(X) \le \sigma$.

Claim 7.8. *If condition (7.14) holds, then* $\max_{X\subseteq V} q'(X) = \max_{X\subseteq V} q(X) \le \sigma$.

Proof. Let X be the set where the maximum is reached for q', and let $\mathcal{F}_1, \mathcal{F}_2$ be tree-compositions of X for which $q'(X) = Q_X(\mathcal{F}_1, \mathcal{F}_2)$. We transform $\mathcal{F}_1$ and $\mathcal{F}_2$ using the following operations until none of them is applicable (see Figure 7.1):

(i) If $Z_1, Z_2 \in \mathcal{F}_1$ are crossing, then replace Z_1, Z_2 by $Z_1 \cap Z_2, Z_1 \cup Z_2$ in $\mathcal{F}_1$.

(ii) If $Z_1, Z_2 \in \mathcal{F}_2$ are crossing, then replace Z_1, Z_2 by $Z_1 \cap Z_2, Z_1 \cup Z_2$ in $\mathcal{F}_2$.

(iii) If $\mathcal{F}_2$ is a partition of some $Z \subseteq V$, and $Z_1 \in \mathcal{F}_1$ and $Z_2 \in \mathcal{F}_2$ are crossing, then replace Z_1 by $Z_1 - Z_2$ in $\mathcal{F}_1$, and replace Z_2 by $Z_2 - Z_1$ in $\mathcal{F}_2$.

(iv) If $\{Z_1, \dots, Z_t\} \subset \mathcal{F}_1$ or $\{Z_1, \dots, Z_t\} \subset \mathcal{F}_2$ is a partition or a co-partition of V, then remove $Z_1, \dots, Z_t$ from that family.

(v) If $\mathcal{F}_2$ is a composition of $Z \subseteq V$ and it contains a subfamily $\{Z_1, \dots, Z_t\}$ $(t \ge 2)$ of pairwise co-disjoint sets such that $\emptyset \ne \cap Z_i \subseteq Z$, then remove $Z_1, \dots, Z_t$ from $\mathcal{F}_2$, and add $V - Z_1, \dots, V - Z_t$ to $\mathcal{F}_1$.

Throughout the process, $\mathcal{F}_1$ and $\mathcal{F}_2$ are compositions. The cardinality of the set for which they are both compositions decreases in operation (iii), while operations (iv) and (v) decrease either $h_X(\mathcal{F}_1)$ or $h_X(\mathcal{F}_2)$. Thus by Lemma 2.4 the process terminates after a finite number of steps. We denote by $\mathcal{F}'_1$ and $\mathcal{F}'_2$ the families obtained at the end. It is easy to see that $\mathcal{F}'_1$ and $\mathcal{F}'_2$ are tree-compositions of some $X' \subseteq X$, $\mathcal{F}'_1 + \mathcal{F}'_2$ is cross-free, and $h_{X'}(\mathcal{F}'_2) \leq 0$ by Claim 2.2. Moreover, $Q_X(\mathcal{F}_1, \mathcal{F}_2)$ does not decrease in any of the steps (in (i)–(iii) this follows from the supermodularity of p, in (iv) it is a consequence of Claim 7.6, and in (v) it can be deduced from the definition of $Q_X(\mathcal{F}_1, \mathcal{F}_2)$). This proves that $\max_{X\subseteq V} q'(X) = \max_{X\subseteq V} q(X)$. □

By Claim 7.8, $\sigma \geq \max_{X\subseteq V} q'(X)$. Thus there exists a vector $m : V \to \mathbb{Z}_+$ with $m(V) = \sigma$ that satisfies (7.1)–(7.3), therefore by Theorem 7.1 there exists a hypergraph H with γ hyperedges of total size σ such that $H_0 + H$ has an orientation covering p, and H satisfies (7.15). This concludes the proof of Theorem 7.5. □

If the requirement function p is monotone decreasing or symmetric, then Theorem 7.5 can be simplified.

Theorem 7.9. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $p : 2^V \to \mathbb{Z}_+$ a monotone decreasing or symmetric non-negative crossing supermodular set function, $\sigma \geq 0$ and $0 \leq \gamma \leq \sigma/2$ integers. There exists a hypergraph H with γ hyperedges of total size σ such that $H_0 + H$ has an orientation covering p if and only if the following hold:*

$$\gamma \geq \sum_{Z\in\mathcal{F}} p(Z) - e_{H_0}(\mathcal{F}) \text{ for every partition } \mathcal{F}, \tag{7.17}$$

$$\sigma \geq \sum_{Z\in\mathcal{F}_1+\mathrm{co}(\mathcal{F}_2)} p(Z) - e_{H_0}(\mathcal{F}_1 + \mathrm{co}(\mathcal{F}_2)) \tag{7.18}$$

whenever $\mathcal{F}_1$ and $\mathcal{F}_2$ are partitions of some $X \subseteq V$ and $\mathcal{F}_1$ is a refinement of $\mathcal{F}_2$. In addition, H can be chosen so that

$$\left\lfloor \frac{\sigma}{\gamma} \right\rfloor \leq |e| \leq \left\lceil \frac{\sigma}{\gamma} \right\rceil \quad \text{for every } e \in \mathcal{E}. \tag{7.19}$$

Proof. It suffices to show that if condition (7.14) is violated for some pair $(\mathcal{F}_1, \mathcal{F}_2)$, then it is also violated by a pair $(\mathcal{F}'_1, \mathcal{F}'_2)$ that has the additional properties that $\mathcal{F}'_1$ is a subpartition, and $Y_2 \not\subset Y_1$ for every $Y_1 \in \mathcal{F}'_1$, $Y_2 \in \mathcal{F}'_2$. Such families can be obtained from $\mathcal{F}_1$ and $\mathcal{F}_2$ by repeating the following operations as long as any of them is possible:

(i) If $\mathcal{F}_1$ is a composition of $Z \subseteq V$, it contains a subfamily $\{W_1, \dots, W_s\}$ $(s \geq 2)$ of pairwise co-disjoint sets such that $W := \cap W_i \subseteq Z$, and $\mathcal{F}_2$ contains a partition $\{Z_1, \dots, Z_t\}$ of W, then remove $W_1, \dots, W_s$ from $\mathcal{F}_1$ and $Z_1, \dots, Z_t$ from $\mathcal{F}_2$.

(ii) If $\mathcal{F}_1$ is a composition of $Z \subseteq V$ and it contains a set $W \subseteq Z$ such that $\mathcal{F}_2$ contains a nontrivial partition $\{Z_1, \dots, Z_t\}$ of W, then replace W in $\mathcal{F}_1$ by the sets $Z_1, \dots, Z_t$, and replace $Z_1, \dots, Z_t$ in $\mathcal{F}_2$ by W.

After a finite number of steps, none of the above operations are applicable; let $(\mathcal{F}_1', \mathcal{F}_2')$ be the pair obtained at that point. Then it follows from Claim 2.2 that $\mathcal{F}_1'$ must be a partition of some $X' \subseteq V$, and operation (ii) guarantees that $\mathcal{F}_1'$ is a refinement of $\mathcal{F}_2'$ if $\mathcal{F}_2' \neq \emptyset$. If (7.17) holds, then the value $Q_X(\mathcal{F}_1, \mathcal{F}_2)$ does not decrease during the above two operations. For (i) this follows from (7.16), since the symmetry or monotone decreasing property of p implies that $Q_W(\{W_1, \dots, W_s\}, \{Z_1, \dots, Z_t\}) \leq Q_\emptyset(\{V - W_1, \dots, V - W_s, Z_1, \dots, Z_t\}, \emptyset) \leq 0$. For (ii), it follows because $p(W) + \sum_{i=1}^t p(V - Z_i) \leq p(V - W) + \sum_{i=1}^t p(Z_i)$, again due to the symmetry or the monotone decreasing property of p. □

Remark. The following example shows that (7.17) itself is not sufficient in Theorem 7.9. Let $V = \{v_1, v_2, v_3, v_4\}$, $E = \{v_1v_2, v_1v_3, v_1v_4\}$. Let $p = 1$ on the sets $\{v_2\}, \{v_3\}, \{v_4\}$ and on their complement; let $p = 0$ on all other sets. If we allow only the addition of graph edges, we need at least 2 of them for a feasible orientation (two edges suffice, since after adding v_2v_3 and v_3v_4 the graph has a strong orientation) but (7.17) requires only $\gamma \geq 1$, since the only deficient partitions are $\{\{v_i\}, V - v_i\}$ $(i = 2, 3, 4)$.

7.3 (k,l)-partition-connectivity augmentation

If $k \geq l$, then the set function p_{kl} defined in (6.3) is non-negative, monotone decreasing, and crossing supermodular, therefore Theorems 7.4 and 7.9 can be applied when it is the requirement function. Since by Theorem 6.12 a hypergraph is (k,l)-partition-connected for $k \geq l$ if and only if it has a (k,l)-edge-connected orientation, we obtain the following:

Corollary 7.10. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $m : V \to \mathbb{Z}_+$ a degree specification, $0 \leq \gamma \leq m(V)/2$ an integer, and $k \geq l$ non-negative integers. There exists a hypergraph H with γ hyperedges such that $H_0 + H$ is (k,l)-partition-connected and $d_H(v) = m(v)$ for all $v \in V$ if and only if the following hold for every nontrivial partition $\mathcal{F}$ of V:*

$$\gamma \geq (|\mathcal{F}| - 1)k + l - e_{H_0}(\mathcal{F}) \ , \tag{7.20}$$

$$\min_i m(V - X_i) \geq (|\mathcal{F}| - 1)k + l - e_{H_0}(\mathcal{F}) \ . \tag{7.21}$$

In addition, H can be chosen so that

$$\left\lfloor \frac{m(V)}{\gamma} \right\rfloor \leq |e| \leq \left\lceil \frac{m(V)}{\gamma} \right\rceil \quad \text{for every } e \in \mathcal{E}. \tag{7.22}$$

□

Note that Corollary 7.10 has some interest even in the very special case when $H_0 = \emptyset$. A result of Edmonds [16] states that a degree-sequence $m_1, \dots, m_n$ is realizable by a k-edge-connected graph if and only if $\sum_{i=1}^n m_i$ is even, and $m_i \ge k$ for every i. Corollary 7.10 implies the following similar result: given $k \ge l$, a degree-sequence $m_1, \dots, m_n$ is realizable by a (k,l)-partition-connected graph if and only if $\sum_{i=1}^n m_i$ is at least $2k(n-1)+2l$, it is even, and $m_i \ge k+l$ for every i.

When $l = 0$, this implies the following tiny result (which is not difficult to prove directly either): If a degree-sequence is realizable by a k-edge-connected graph with at least $k(n-1)$ edges, then it is also realizable by a graph containing k edge-disjoint spanning trees.

On minimum cardinality augmentation, Theorem 7.9 gives the following:

Corollary 7.11. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $\sigma \ge 0$, $0 \le \gamma \le \sigma$, and $k \ge l$ non-negative integers. There is a hypergraph H with γ hyperedges of total size σ such that $H_0 + H$ is (k,l)-partition-connected if and only if the following two conditions are met:*

(i) $\gamma \ge (|\mathcal{F}|-1)k + l - e_{H_0}(\mathcal{F})$ for every nontrivial partition $\mathcal{F}$,

(ii) $\sigma \ge |\mathcal{F}_1|k + |\mathcal{F}_2|l - e_{H_0}(\mathcal{F}_1 + \mathrm{co}(\mathcal{F}_2))$ whenever $\mathcal{F}_1$ and $\mathcal{F}_2$ are partitions of some $\emptyset \ne X \subset V$ and $\mathcal{F}_1$ is a refinement of $\mathcal{F}_2$.

In addition, H can be chosen so that

$$\left\lfloor \frac{\sigma}{\gamma} \right\rfloor \le |e| \le \left\lceil \frac{\sigma}{\gamma} \right\rceil \quad \textit{for every } e \in \mathcal{E}. \tag{7.23}$$

□

Remark. The graph on Figure 7.2 shows that the second condition in Corollary 7.11 cannot be simplified. We need to add at least 2 edges to the graph to have a $(3,1)$-edge-connected orientation from root s, but the simplest evidence for this is the family indicated on the figure (consisting of the round sets and the complements of the square sets), whose deficiency is 3, while a new edge can enter at most 2 sets. The figure on the right shows that the addition of 2 edges is sufficient (to see that the digraph is $(3,1)$-edge-connected, observe that it contains 3 edge-disjoint out-arborescences from s, and also an in-arborescence to s).

Let $\nu \ge 2$ be an integer. Because of (7.23), Corollary 7.11 gives a characterization for minimum cardinality augmentation with ν-hyperedges. A simple observation shows that a characterization of minimum total size augmentation can also be derived. It is easy to see that we get the best bound if $\gamma = \lfloor \sigma/2 \rfloor$:

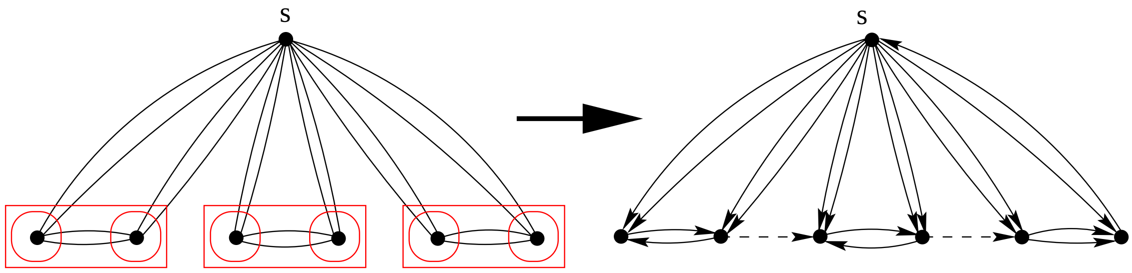

Figure 7.2: $(3,1)$-partition-connectivity augmentation by graph edges. Condition (ii) in Corollary 7.11 is tight on the indicated family (the non-rounded rectangles represent their complement). The picture on the right shows a $(3,1)$-edge-connected orientation of the augmented graph.

Corollary 7.12. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $\sigma \geq 0$ and $k \geq l$ non-negative integers. There is a hypergraph H of total size σ such that $H_0 + H$ is (k,l)-partition-connected if and only if the following two conditions are met:*

(i) $\frac{\sigma}{2} \geq (|\mathcal{F}| - 1)k + l - e_{H_0}(\mathcal{F})$ *for every nontrivial partition* $\mathcal{F}$,

(ii) $\sigma \geq |\mathcal{F}_1|k + |\mathcal{F}_2|l - e_{H_0}(\mathcal{F}_1 + \mathrm{co}(\mathcal{F}_2))$ *whenever* $\mathcal{F}_1$ *and* $\mathcal{F}_2$ *are partitions of some* $X \subset V$ *and* $\mathcal{F}_1$ *is a refinement of* $\mathcal{F}_2$.

In addition, there is an optimal H that contains only graph edges and at most one 3-hyperedge.

7.4 More general requirement functions

In this section we illustrate what happens if instead of considering only non-negative crossing supermodular requirement functions, we allow the broader class of positively crossing supermodular functions. We were only able to solve the degree-specified problem in this case, and the characterizations are not elegant. However, the result provides an extension of combined augmentation and orientation to mixed graphs and mixed hypergraphs, so it may be of some interest.

7.4.1 Mixed hypergraphs

The orientation of mixed graphs was the main motivation behind Theorem 6.5 of Frank [25], which extended Theorem 6.3. We intend to generalize Theorem 7.1 in a similar way. A *mixed hypergraph* $M = (V; \mathcal{E}, \mathcal{A})$ contains a set $\mathcal{E}$ of hyperedges and a set $\mathcal{A}$ of hyperarcs;

this notion already appeared in Theorem 6.13 which was an extension of the results of Khanna, Naor, and Shepherd [50] on network design with orientation constraints.

The problem of finding an orientation of a mixed hypergraph $M = (V; \mathcal{E}, \mathcal{A})$ that covers a set function p is equivalent to the problem of finding an orientation of the hypergraph $H = (V, \mathcal{E})$ that covers the set function $p'(X) = (p(X) - \varrho_{\mathcal{A}}(X))^+$. While p' is not necessarily crossing supermodular if p is such, it certainly is positively crossing supermodular. This motivates the study of combined augmentation and orientation problems for positively crossing supermodular requirement functions. Minimum cardinality augmentation (even with graph edges) remains an open problem, but we present a solution for the degree-specified problem.

Remark. In Theorem 6.5, crossing supermodular set functions with possible negative values were considered. The following example shows that the positively crossing supermodular case is more general, i.e. not every positively crossing supermodular set function p can be made crossing supermodular by decreasing the value of p on some of the sets where it is 0.

Let X_1, X_2, X_3 be three subsets of a ground set V, in general position. Let $p(X_i) = 1$, $p(X_i \cup X_j) = 2$ $(i \neq j)$, $p(X_1 \cup X_2 \cup X_3) = 4$, and $p(X) = 0$ on the remaining sets; this is a positively crossing supermodular function. The value of $p(X_1 \cap X_2)$ cannot be decreased since

$$p(X_1 \cap X_2) \geq p(X_1) + p(X_2) - p(X_1 \cup X_2) = 0 \ .$$

Therefore it is impossible to correctly modify p so as to satisfy

$$p(X_1 \cap X_2) \leq p(X_1 \cap X_2 \cap X_3) + p(X_1 \cap X_2 \cup X_3) - p(X_3) \leq -1 \ .$$

7.4.2 Degree specified augmentation

As in Theorem 6.10, our result gives characterizations for both the case when the added hypergraph must be uniform and when there is no such requirement. When the degree-specification is 0 on every node and H_0 is a graph, we obtain Theorem 6.5. Since the characterizations in that theorem already included tree-compositions, we cannot expect anything less complicated here.

Theorem 7.13. *Let $H_0 = (V, \mathcal{E}_0)$ be a hypergraph, $p : 2^V \to \mathbb{Z}_+$ a positively crossing supermodular set function, $m : V \to \mathbb{Z}_+$ a degree specification, and $0 \leq \gamma \leq m(V)/2$ an integer. We define the set function*

$$p^m(X) := p(X) + (\gamma - m(V - X))^+ \ .$$

There exists a hypergraph $H = (V, \mathcal{E})$ with γ hyperedges such that H_0+H has an orientation covering p and $d_H(v) = m(v)$ for every $v \in V$ if and only if

$$\sum_{Z\in\mathcal{F}} p^m(Z) + (\gamma - m(X))^+ \leq e_{H_0}(\mathcal{F}) + (h_X(\mathcal{F}) + 1)\gamma \tag{7.24}$$

for every $X \subseteq V$ and every tree-composition $\mathcal{F}$ of X. In addition, H can be chosen so that

$$\left\lfloor \frac{m(V)}{\gamma} \right\rfloor \leq |e| \leq \left\lceil \frac{m(V)}{\gamma} \right\rceil \quad \textit{for every } e \in \mathcal{E}. \tag{7.25}$$

Proof. The necessity follows from the fact that if $\mathcal{F}$ is a tree-composition of X, then $\mathcal{F}' := \mathcal{F} + \{V - X\}$ is a regular family, $\sum_{Z\in\mathcal{F}'} \varrho_{\vec{H}_0}(Z) \leq e_{H_0}(\mathcal{F}')$ for any orientation $\vec{H}_0$ of H_0, and $\sum_{Z\in\mathcal{F}'} \varrho_{\vec{H}}(Z) \leq h_\emptyset(\mathcal{F}')\gamma - \sum_{Z\in\mathcal{F}'} (\gamma - m(V - Z))^+$ for any orientation $\vec{H}$ of a hypergraph H satisfying the degree specification.

The sufficiency can be proved in essentially the same way as in the proof of Theorem 7.1. We define H_0' similarly, and for $X \subseteq V$, let

$$p_1(X) := i_{H_0}(X) + (\gamma - m(V - X))^+ ,$$

$$p_2(X) := p(X) + i_{H_0}(X) + (\gamma - m(V - X))^+ .$$

In this case Lemma 6.10 implies that an orientation of H_0' satisfying (7.5)–(7.7) exists if and only if the polyhedron

$$\{x : V \to \mathbb{Q} :\ x(V) = p_1(V),\ x(Z) \geq p_1(Z)\ \forall Z \subseteq V,\ x(Z) \geq p_2(Z)\ \forall Z \subseteq V\}$$

has an integral point.

Claim 7.14. *The set function p_1 is fully supermodular, and the set function p_2 is supermodular on the crossing pairs (X, Y) for which $p_1(X) < p_2(X)$ and $p_1(Y) < p_2(Y)$.*

Proof. The set function p_1 is the sum of two fully supermodular functions (see the proof of Claim 7.2), so it is fully supermodular. Since p is positively crossing supermodular, p_2 is supermodular on the crossing pairs (X, Y) for which $p(X) > 0$ and $p(Y) > 0$, and these are exactly the crossing pairs for which $p_1(X) < p_2(X)$ and $p_1(Y) < p_2(Y)$. □

Theorem 2.24 implies that an orientation of H_0' satisfying (7.5)–(7.7) exists if and only if

$$p_1(V - X) + \sum_{Z\in\mathcal{F}} p_2(Z) \leq (h_X(\mathcal{F}) + 1)p_1(V)$$

for every $X \subseteq V$ and every tree-composition $\mathcal{F}$ of X. Using (1.12) and the fact that $e_{H_0}(\mathcal{F}) = e_{H_0}(\mathcal{F} + \{V - X\})$, this is equivalent to the condition of the theorem.

From here we can follow the line of the proof of Theorem 7.1. Let $\vec{H}'_0$ be the orientation of H'_0 satisfying (7.5)–(7.7), and let $\vec{H}_0$ denote the directed hypergraph obtained from $\vec{H}'_0$ by deleting the node z. Let $m_i(v)$ be the multiplicity of the edge zv in $\vec{H}'_0$, and $m_o(v)$ the multiplicity of the edge vz in $\vec{H}'_0$. We define the set function $q(X) = (p(X) - \varrho_{\vec{H}_0}(X))^+$; then q is positively crossing supermodular. As in the proof of Theorem 7.1, we can apply Theorem 4.7 (with the m_i, m_o and q defined above) to obtain a directed hypergraph D whose underlying undirected hypergraph H is a good augmentation of H_0. Theorem 4.7 also ensures that (7.25) can be satisfied. This concludes the proof of Theorem 7.13. $\square$

Chapter 8

Concluding remarks

In these paragraphs we give a brief account of the open problems related to the various topics discussed in the thesis. Most of these questions were mentioned in previous chapters; they are collected here in order to indicate some possible directions for future research.

One of the new results described in the thesis is Corollary 3.30 on k-edge-connectivity augmentation of hypergraphs by the addition of uniform hyperedges. This result followed from Theorems 3.14 and 3.29, which addressed the more general problem of covering symmetric crossing supermodular set functions by uniform hypergraphs. The proofs included an extension of the well-known splitting-off technique to hypergraphs. While the existence of a special splitting-off sequence was shown, it would be desirable to obtain a more complete structural description of the feasible splitting-off operations. This might help in dealing with problems with extra constraints, like augmentation with partition constraints, or simultaneous augmentation of two hypergraphs.

For directed hypergraphs, the splitting-off method used in the proof of Theorem 4.7 is essentially the same as the one described by Berg, Jackson, and Jordán in [7]. Theorem 4.7 shows that if $m_i(V) \leq m_o(V)$ then there is always a complete splitting-off preserving k-edge-connectivity in V. However, this is not necessarily true if only digraph edges are allowed as new edges. Therefore in this case it is natural to ask what is the maximum length of a feasible splitting-off sequence. Berg, Jackson, and Jordán conjectured the following:

Conjecture 8.1. *Let $D = (V + s, \mathcal{A})$ be a directed hypergraph that is k-edge-connected in V, and suppose that s is incident only with digraph edges and $\varrho(s) \geq \delta(s)$. Then D does not have a sequence of γ feasible $(1,1)$-splittings at s if and only if there is a family $\mathcal{F} = \{X_1, X_2, \dots, X_t\}$ of pairwise co-disjoint subsets of V for some $2 \leq t \leq 2\gamma + 1$, such that*

$$\sum_{i=1}^{t} \varrho(X_i) \leq tk + (t-1)\gamma - \mu - 1,$$

where $\mu = \min\{\gamma, |\{sv \in \mathcal{A} : v \notin \cap_{i=1}^{t} X_i\}|\}$.

Let us return to undirected hypergraphs. We have shown in Chapter 5 how matroid-theoretic tools applied on hypergraphic matroids can be used to prove results like Theorem 5.7 on the decomposition of a hypergraph into k partition-connected sub-hypergraphs. This result in turn easily implies Theorem 5.12 which gives a sufficient condition for the existence of k edge-disjoint Steiner trees for W when $V - W$ is stable. Kriesell [53] conjectures that $2k$-edge-connectivity in W is sufficient, even in the case when $V - W$ is not necessarily stable. To date, no constant c is known such that ck-edge-connectivity in W guarantees the existence of k edge-disjoint Steiner trees.

Another open problem for graphs that is related to partition-connectivity is the constructive characterization of (k, l)-partition-connected graphs, i.e. Conjecture 6.7. This conjecture would follow from Conjecture 6.6 on the constructive characterization of (k, l)-edge-connected digraphs. Since by Theorem 6.11 there is a similar relation between (k, l)-partition-connectivity and (k, l)-edge-connectivity for hypergraphs, we might obtain similar constructive characterizations for hypergraphs as well.

One of the exciting questions in graph connectivity orientation is the solvability of orientation problems with parity constraints. Known results were cited in Chapter 6; a basic open problem is the characterization of graphs which have a strongly connected orientation where the in-degree of every node is odd. Of course, similar questions can be asked for hypergraphs.

A problem where the answer is known for graphs but not for hypergraphs is local edge-connectivity orientation with symmetric demands. For example, no characterization is known for hypergraphs which have a k-edge-connected orientation that is at the same time l-edge-connected in a given subset W $(l > k)$. For graphs, this problem is solved by Theorem 6.4 of Nash-Williams, whose proof makes use of the basic properties of Eulerian graphs; so it might be interesting to find an appropriate definition of "Eulerian hypergraphs" that is useful in this context.

One can also extend orientation problems to mixed graphs and hypergraphs. Instances where we obtain tractable problems are the framework of directed network design with orientation constraints, introduced by Khanna, Naor, and Shepherd (an extended hypergraph version is described in Theorem 6.14), and Theorem 6.5 of Frank. The latter can be seen as a special case of the combined augmentation and orientation problem solved in Theorem 7.13. This theorem gives a characterization for a degree-specified augmentation problem; we do not yet know how to solve the corresponding minimum cardinality problem.

Bibliography

[1] J. Bang-Jensen, H. Gabow, T. Jordán, and Z. Szigeti, *Edge-connectivity augmentation with partition constraints*, SIAM J. Discrete Mathematics 12 No. 2 (1999), 160–207.

[2] J. Bang-Jensen, B. Jackson, *Augmenting hypergraphs by edges of size two*, in: Connectivity Augmentation of Networks: Structures and Algorithms, Mathematical Programming (ed. A. Frank) Ser. B, Vol. 84 No. 3 (1999), 467–481.

[3] J. Bang-Jensen, T. Jordán, *Edge-connectivity augmentation preserving simplicity*, SIAM J. Discrete Math. Vol. 11 No. 4 (1998), 603–623.

[4] J. Bang-Jensen, A. Frank, B. Jackson, *Preserving and increasing local edge-connectivity in mixed graphs*, SIAM J. Discrete Math. 8 No. 2 (1995), 155–178.

[5] J. Bang-Jensen, S. Thomassé, *Decompositions and orientations of hypergraphs*, IMADA Preprint 2001-10.

[6] A. A. Benczúr, D.R. Karger, *Quadratic-time algorithms for edge-connectivity augmentation and splitting off*, J. Alg 37 (2000), 2–36.

[7] A. Berg, B. Jackson, T. Jordán, *Edge splitting and connectivity augmentation in directed hypergraphs*, Discrete Mathematics, to appear

[8] A. Berg, B. Jackson, T. Jordán, *Highly edge-connected detachments of graphs and digraphs*, J. Graph Theory Vol. 43, No. 3 (2003), 67–77.

[9] A. Benczúr, A. Frank, *Covering symmetric supermodular functions by graphs*, in: Connectivity Augmentation of Networks: Structures and Algorithms, Mathematical Programming (ed. A. Frank), Ser. B, Vol. 84, No. 3 (1999), 483–503.

[10] F. Boesch, R. Tindell, *Robbins's theorem for mixed multigraphs*, Am. Math. Monthly 87 (1980), 716–719.

[11] E. Cheng, *Edge-augmentation of hypergraphs*, in: Connectivity Augmentation of Networks: Structures and Algorithms, Mathematical Programming (ed. A. Frank), Ser. B, Vol. 84, No. 3 (1999), 443–465.

[12] B. Cosh, B. Jackson, Z. Király, *Local connectivity augmentation is NP-complete*, appeared in: B. Cosh: Vertex splitting and connectivity augmentation in hypergraphs, Ph.D. thesis, Goldsmiths College, University of London (2000).

[13] J. Edmonds, *Minimum partition of a matroid into independent sets*, J. Res. Nat. Bur. Standards Sect. 869 (1965), 67–72.

[14] J. Edmonds, *Edge-disjoint branchings*, in: B. Rustin, ed., Combinatorial Algorithms (Acad. Press, New York, 1973) 91–96.

[15] J. Edmonds, *Submodular functions, matroids, and certain polyhedra*, in: R. Guy, H. Hanani, N. Sauer, and J. Schönheim (eds.): Combinatorial Structures and their Applications. Gordon and Breach, New York (1970), 69–87.

[16] J. Edmonds, *Existence of k-edge-connected ordinary graphs with prescribed degrees*, J. Res. Nat. Bur. Standards Sect. B 68B (1964), 73–74.

[17] J. Edmonds, R. Giles, *A min-max relation for submodular functions on graphs*, Ann. Discrete Math. 1 (1977), 185–204.

[18] L.P. Eswaran, R.E. Tarjan, *Augmentation problems*, SIAM J. Computing 5 No. 4 (1976), 653–665.

[19] B. Fleiner, *Detachment of vertices preserving edge-connectivity*, SIAM J. Discrete Math., to appear.

[20] T. Fleiner, T. Jordán, *Coverings and structure of crossing families*, in: Connectivity Augmentation of Networks: Structures and Algorithms, Mathematical Programming (ed. A. Frank), Ser. B, Vol. 84, No. 3 (1999), 505–518.

[21] L.R. Ford, D.R. Fulkerson, *Flows in Networks*, Princeton Univ. Press, Princeton NJ., 1962.

[22] A. Frank, *On the orientation of graphs*, J. Combinatorial Theory B 28 No. 3 (1980), 251–261.

[23] A. Frank, *Finding feasible vectors of Edmonds-Giles polyhedra*, Journal of Combinatorial Theory Series B 36 (1984), 221–239.

[24] A. Frank, *Connectivity augmentation problems in network design*, in: J.R. Birge, K.G. Murty (eds.): Mathematical Programming: State of the Art. Univ. of Michigan (1994), 34–63.

[25] A. Frank, *Orientations of graphs and submodular flows*, Congressus Numerantium 113 (1996), 111–142.

[26] A. Frank, *Matroids and submodular functions*, in: Annotated Bibliographies in Combinatorial Optimization (M. Dell'Amico, F. Maffioli, S. Martello, eds.), 1997, John Wiley, 65–80.

[27] A. Frank, *On disjoint trees and arborescences*, in: Algebraic Methods in Graph Theory, Colloquia Mathematica Soc. J. Bolyai 25, North Holland (1978), 159–169.

[28] A. Frank, *Kernel systems of directed graphs*, Acta Sci. Math. (Szeged) 41, 1-2 (1979), 63–76.

[29] A. Frank, *Augmenting graphs to meet edge-connectivity requirements*, SIAM J. Discrete Math. 5 No. 1 (1992) 22–53.

[30] A. Frank, *On a theorem of Mader*, Annals of Discrete Mathematics 101 (1992) 49–57.

[31] A. Frank, Z. Király, *Graph orientation with edge-connection and parity constraints*, Combinatorica 22 (2002), 47–70.

[32] A. Frank, É. Tardos, *Generalized polymatroids and submodular flows*, Mathematical Programming, Ser. B, Vol. 42 (1988), 489–563.

[33] A. Frank, T. Jordán, Z. Szigeti, *An orientation theorem with parity conditions*, Discrete Applied Mathematics 115 (2001), 37–45.

[34] A. Frank, L. Szegő, *Constructive characterizations for packing and covering with trees*, Discrete Applied Mathematics, to appear

[35] A. Frank, T. Király, *Combined connectivity augmentation and orientation problems*, Discrete Applied Mathematics, to appear

[36] A. Frank, T. Király, Z. Király, *On the orientation of graphs and hypergraphs*, Discrete Applied Mathemathics, to appear

[37] A. Frank, T. Király, M. Kriesell, *On decomposing a hypergraph into k connected sub-hypergraphs*, Discrete Applied Mathematics, to appear

[38] S. Fujishige, *Structures of polyhedra determined by submodular functions on crossing families*, Mathematical Programming 29 (1984), 125–141.

[39] H.N. Gabow, *Efficient splitting off algorithms for graphs*, in: Proceedings of the 26th ACM Symposium on Theory of Computing, 696–705.

[40] G. Gallo, G. Longo, S. Nguyen and S. Pallottino, *Directed hypergraphs and applications*, Discrete Appl. Math. 40 (1993), 177–201.

[41] G. Gallo, M.G. Scutellà, *Directed Hypergraphs as a Modelling Paradigm*, Rivista AMASES 21 (1998), 97–123.

[42] R.E. Gomory, T.C. Hu, *Multi-terminal network flows*, SIAM J. Appl. Math 9 (1961) 551-570.

[43] M. Grötschel, L. Lovász, A. Schrijver, *The ellipsoid algorithm and its consequences in combinatorial optimization*, Combinatorica 1 (1981), 499–513.

[44] S.L. Hakimi, *On the degrees of the vertices of a directed graph*, Journal of the Franklin Institute 279 No. 4 (1965), 290–308.

[45] S. Iwata, L. Fleischer, S. Fujishige, *A combinatorial, strongly polynomial-time algorithm for minimizing submodular functions*, J. ACM 48 (2001), 761–777.

[46] K. Jain, M. Mahdian, M.R. Salavatipour, *Packing Steiner trees*, SODA 2003.

[47] T. Jordán, *Two NP-complete augmentation problems*, IMADA, Odense Universitet Preprints No. 8/1997.

[48] T. Jordán, *Constrained edge-splitting problems*, SIAM J. Discrete Math., to appear.

[49] T. Jordán, Z. Szigeti, *Detachments preserving local edge-connectivity of graphs*, SIAM J. Discrete Math., to appear.

[50] S. Khanna, J. Naor and F. B. Shepherd, *Directed network design with orientation constraints*, in: Proceedings of the Eleventh Annual ACM–SIAM Symposium on Discrete Algorithms (2000), 663–671.

[51] T. Király, *Covering symmetric supermodular functions by uniform hypergraphs*, submitted (Journal of Combinatorial Theory Series B).

[52] T. Király, M. Makai, *A note on hypergraph connectivity augmentation*, EGRES Technical Report no. 2002-11.

[53] M. Kriesell, *Local spanning trees in graphs and hypergraph decomposition with respect to edge-connectivity*, Report 257 University of Hannover (1999), Submitted for publication.

[54] M. Lorea, *Hypergraphes et matroides*, Cahiers Centre Etud. Rech. Oper. 17 (1975), 289–291.

[55] L. Lovász, *Combinatorial Problems and Exercises*, North-Holland (1979).

[56] L. Lovász, *A generalization of Kőnig's theorem*, Acta. Math. Acad. Sci. Hungar. 21 (1970), 443–446.

[57] L. Lovász, *Flats in matroids and geometric graphs*, in: Combinatorial Surveys, Proceedings of the 6th British Combinatorial Conference, Academic Press, London, 1977, 45–86.

[58] W. Mader, *A reduction method for edge-connectivity in graphs*, Ann. Discrete Math. 3 (1978), 145–164.

[59] W. Mader, *Konstruktion aller n-fach kantenzusammenhängenden Digraphen*, Europ. J. Combinatorics 3 (1982), 63–67.

[60] H. Nagamochi, P. Eades, *Edge-splitting and edge-connectivity augmentation in planar graphs*, IPCO IV Springer LNCS 1412, R.E. Bixby et al, eds. (1998), 96–111.

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

[62] C.St.J.A Nash-Williams, *Well-balanced orientations of finite graphs and unobtrusive odd-vertex pairings*, in: Recent Progress in Combinatorics (ed. W.T. Tutte), Academic Press (1969), 133–149.

[63] C.St.J.A. Nash-Williams, *Decomposition of finite graphs into forests*, J. London Math. Soc. 39 (1964), 12.

[64] J. Plesnik, *Minimum block containing a given graph*, Archiv der mathematik Vol. XXVII (1976), 668–672.

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

[66] A. Schrijver, *Total dual integrality from directed graphs, crossing families and sub- and supermodular functions*, in: W. R. Pulleyblank, ed., Progress in Combinatorial Optimization, Academic Press (1984), 315–362.

[67] A. Schrijver, *A combinatorial algorithm minimizing submodular functions in strongly polynomial time*, J. Combin. Theory Ser. B 80 (2000), 346–355.

[68] A. Schrijver, *Combinatorial Optimization - Polyhedra and Efficiency*, Springer-Verlag, Berlin (2003).

[69] Z. Szigeti, *Hypergraph connectivity augmentation*, in: Connectivity Augmentation of Networks: Structures and Algorithms, Mathematical Programming (ed. A. Frank), Ser. B, Vol. 84, No. 3 (1999), 519–527.

[70] W.T. Tutte, *On the problem of decomposing a graph into n connected factors*, J. London Math. Soc. 36 (1961), 221–230.

[71] T. Watanabe, A. Nakamura, *Edge-connectivity augmentation problems*, Computer and System Sciences 35 No. 1 (1987), 96–144.

Index

Notation

$\mathbb{Z}_+$	the set of non-negative integers
$(x)^+$	for $x \in \mathbb{R}$: $\max\{x, 0\}$
$x > 0$	for $x \in \mathbb{R}^n$: $x \geq 0$, $x \neq 0$
$x \gg 0$	for $x \in \mathbb{R}^n$: $x_i > 0$ $(i = 1, \ldots, n)$
$X \subseteq V$	X is a subset of V (possibly $X = V$)
$X \subset V$	X is a proper subset of V
$X - Y$	for $X, Y \subseteq V$: $\{v \in V : \ v \in X, \ v \notin Y\}$
χ_X	for $X \subseteq V$: the characteristic function of the set X
$m(X)$	for $X \subseteq V$ and $m : V \to \mathbb{R}$: $\sum_{v \in X} m(v)$
$d_H(v)$	for a hypergraph $H = (V, \mathcal{E})$ and $v \in V$: the degree of v in H
$\varrho_D(v)$	for a dir. hypergraph $D = (V, \mathcal{A})$: the in-degree of v in D
$\delta_D(v)$	the out-degree of v in D
$\Delta_H(X)$	$\{e \in \mathcal{E} : \ e$ enters $X\}$
$\Delta_D^-(X)$	$\{a \in \mathcal{A} : \ a$ enters $X\}$
$\Delta_D^+(X)$	$\{a \in \mathcal{A} : \ a$ enters $V - X\}$
$d_H(X)$	$\lvert\Delta_H(X)\rvert$
$d_{\mathcal{E}}(X)$	$d_H(X)$ for $H = (V, \mathcal{E})$
$\varrho_D(X)$	$\lvert\Delta_D^-(X)\rvert$
$\delta_D(X)$	$\lvert\Delta_D^+(X)\rvert$
$i_H(X)$	the number of hyperedges of H induced by X
$\varrho_x(Z)$	for $D = (V, \mathcal{A})$ and $x : \mathcal{A} \to \mathbb{R}$: $\sum_{a \in \Delta_D^-(Z)} x(a)$
$\phi_x(Z)$	$\delta_x(Z) - \varrho_x(Z)$
$\overline{s}t$-set	a set $X \subseteq V$ for which $s \notin X$ and $t \in X$
$\lambda_H(s, t)$	local edge-connectivity between s and t in H
$\lambda_D(s, t)$	local edge-connectivity from s to t in D
$\chi_e(v)$	for $v \in V$ and a hyperedge e: the multiplicity of v in e
$\lvert e \cap X\rvert$	for $X \subseteq V$ and a hyperedge e: $\chi_e(X)$
$\cup(\mathcal{E})$	for $H = (V, \mathcal{E})$: $\{v \in V : \ \chi_e(v) > 0$ for some $e \in \mathcal{E}\}$

$h(a)$	the head node of the hyperarc a
$t(a)$	the multiset of tail nodes of the hyperarc a
$\mathrm{co}(\mathcal{F})$	$\{V - X : \ X \in \mathcal{F}\}$ (with multiplicity)
$\chi_\mathcal{F}$	the characteristic set function of the family $\mathcal{F}$
$\mathcal{F}_1 + \mathcal{F}_2$	family with characteristic function $\chi_{\mathcal{F}_1} + \chi_{\mathcal{F}_2}$
$e_H(\mathcal{F})$	$\sum_{e\in\mathcal{E}} \max_{u\in e} \lvert\{X \in \mathcal{F} : \ u \in X, \ e \not\subseteq X\}\rvert$
$h_X(\mathcal{F})$	for $X \subseteq V$ and a composition $\mathcal{F}$ of X: $\sum_{Z\in\mathcal{F}} \chi_Z(v) - \chi_X(v)$ for an arbitrary $v \in V$
r_M	for a matroid M: the rank function of M
$p^\wedge(X)$	$\max\{\sum_{Z\in\mathcal{F}} p(Z) : \mathcal{F}$ is a partition of $X\}$
$p^\uparrow(X)$	$\max\{\sum_{Z\in\mathcal{F}} p(Z) - h_X(\mathcal{F})p(V) : \mathcal{F}$ is a tree-composition of $X\}$
$p^*(X)$	$\max\{\sum_{Z\in\mathcal{F}} p(Z) - h_X(\mathcal{F})p(V) : \mathcal{F}$ is a composition of $X\}$
$p^*(X, Y)$	if $p(V) = 0$: $\max\{\sum_{Z\in\mathcal{F}} p(Z) : \mathcal{F}$ is an (X, Y)-composition$\}$
$C(p)$	$\{x : V \to \mathbb{Q} : \ x(Y) \geq p(Y) \ \forall Y \subseteq V\}$
$B(p)$	$\{x : V \to \mathbb{Q} : \ x(V) = p(V); \ x(Y) \geq p(Y) \ \forall Y \subseteq V\}$

Abstract

The objective of the thesis is to discuss edge-connectivity and related connectivity concepts in the context of undirected and directed hypergraphs. In particular, we focus on k-edge-connectivity and (k,l)-partition-connectivity of hypergraphs, and (k,l)-edge-connectivity of directed hypergraphs. A strong emphasis is placed on the role of submodularity in the structural aspects of these problems.

One area that is discussed extensively is connectivity augmentation. A min-max theorem is given on the minimum number of ν-hyperedges that have to be added to an initial hypergraph to make it k-edge-connected. Analogously, we prove a formula on the minimum number of $(r,1)$-hyperarcs whose addition makes an initial directed hypergraph (k,l)-edge-connected. These problems (and most others in the thesis) are studied in the general framework of covering supermodular set functions.

We show that matroid techniques can be used in the description of (k,l)-partition-connected hypergraphs. This notion also leads to connectivity orientation problems for hypergraphs, and with these tools we prove characterizations of (k,l)-partition-connectivity and (k,l)-edge-connected orientations. An application concerning edge-disjoint Steiner trees is also given, as well as some new results on directed network design with orientation constraints.

The thesis is concluded with the study of a new class of connectivity augmentation problems, in which the aim is to add hyperedges to an undirected (or mixed) hypergraph such that the resulting hypergraph has an orientation with specified connectivity properties. A special case of the described results is a solution of the (k,l)-partition-connectivity augmentation problem.

The above results are based on the papers [35], [36], [37], [51], and [52]. The thesis also includes a new characterization of set functions defining base polyhedra.

Összefoglaló

A disszertáció irányítatlan és irányított hipergráfok vonatkozásában foglalkozik különböző élösszefüggőségi problémákkal, és ezekhez kapcsolódó különféle összefüggőségi tulajdonságokkal. Részletesen tárgyaljuk például irányítatlan hipergráfok k-élösszefüggőségét és (k, l)-partíció-összefüggőségét, valamint irányított hipergráfok (k, l)-élösszefüggőségét. Kiemelten vizsgáljuk a szubmodularitás szerepét a feladatok strukturális leírásában.

A dolgozatban tárgyalt egyik nagy terület az élösszefüggőség-növelési problémák témaköre. Ezzel kapcsolatban bebizonyítunk egy minimax tételt azon uniform hiperélek minimális számáról, melyek hozzáadásával egy kiindulási hipergráf k-élösszefüggővé tehető. Irányított hipergráfok esetében pedig a (k, l)-élösszefüggővé növeléshez szükséges minimális számú uniform irányított hiperélt adjuk meg. Ezek a problémák, hasonlóan a disszertációban szereplő eredmények zöméhez, belehelyezhetők egy általánosabb keretbe, és leírhatók szupermoduláris halmazfüggvények fedésére vonatkozó feladatként.

Hipergráfok (k, l)-partíció-összefüggőségéről megmutatjuk, hogy a tulajdonság leírható matroidelméleti módszerekkel, és ismertetünk egy Steiner-fákra vonatkozó alkalmazást. Ugyanilyen tulajdonságok játszanak fontos szerepet hipergráfok irányítási problémáiban, és ennek segítségével megadható a (k, l)-partíció-összefüggő hipegráfok egy hatékony jellemzése.

A disszertáció utolsó részében élösszefüggőség-növelési problémák egy újszerű osztályát vizsgáljuk. A cél itt egy irányítatlan vagy vegyes hipergáf növelése oly módon, hogy a kapott hipergráfnak legyen előírt összefüggőségi feltételeket kielégítő irányítása. A leírt eredmények speciális esetének tekinthető a (k, l)-partíció-összefüggőség növelési probléma megoldása.

A fent említett eredmények alapjául a [35], [36], [37], [51], és [52] dolgozatok szolgálnak. A disszertáció ezeken kívül tartalmazza bázis-poliédert definiáló halmazfüggvények egy újszerű jellemzését is.