

EÖTVÖS LORÁND UNIVERSITY
INSTITUTE OF MATHEMATICS



Ph.D. thesis

Connectivity augmentation algorithms

László Végh

Doctoral School: Mathematics

Director: Miklós Laczkovich, member of the Hungarian Academy of Sciences

Doctoral Program: Applied Mathematics

Director: György Michaletzky, Professor, Doctor of Sciences

Supervisor: András Frank

Professor, Doctor of Sciences

Department of Operations Research, Eötvös Loránd University and
MTA-ELTE Egerváry Research Group on Combinatorial Optimization

January 2010

Contents

1	Introduction	1
1.1	The Frank-Jordán Theorem and node-connectivity augmentation	3
1.2	Previous algorithmic results on connectivity augmentation	9
1.3	Undirected edge-connectivity augmentation	11
1.4	Constructive characterizations	15
1.5	Overview of the main results	21
2	Augmenting directed node-connectivity by one	35
2.1	The Dual Oracle	35
2.2	Algorithmic Proof of Theorem 1.6	39
2.3	Further remarks	42
2.4	Implementation via bipartite matching	43
3	Undirected node-connectivity augmentation	47
3.1	Preliminaries	47
3.2	The proof of Theorem 1.37	52
3.3	The Algorithm	57
3.4	Further remarks	61
3.5	Implementation via bipartite matching	65
4	General directed node-connectivity augmentation	69
4.1	The algorithm	70
4.2	Application for directed connectivity augmentation	81
4.3	Further remarks	84
5	Local edge-connectivity augmentation	85
5.1	Coverings without partition constrains	85
5.2	Basic results on partition-constrained local edge-connectivity augmentation	91
5.3	Towards proving the conjectures	99
5.4	Further remarks	112

6	Constructive characterization of (k, ℓ)-edge-connected digraphs	115
6.1	Basic concepts and the proof of Theorem 1.47	115
6.2	Proof of Theorem 6.7	119
6.3	Splitting off	121
6.4	Lemmas	124
6.5	Further remarks	129
	Bibliography	138

Notation

Undirected graphs

$G = (V, E)$	An undirected graph G on node set V with edge set E .
$G = (S, T; E)$	A bipartite graph with colour classes S and T and edge set E .
V^2	The set of all edges on node set V .
$d_G(X)$	The number of edges in G incident to node set X .
$d_G(X, Y)$	The number of edges in G between $X - Y$ and $Y - X$.
$i_G(X)$	The number of edges with both endnodes in X .
$\overline{d}_G(X, Y)$	The number of edges in G between $X \cap Y$ and $V - (X \cup Y)$.
$N_G(X)$	The set of neighbours of node set X .
X^*	$= V - (X \cup N_G(X))$ for node set X .
$\Gamma_G(X)$	The set of neighbours of $X \subseteq S$ or $X \subseteq T$ in a bipartite graph.
$I_G(X)$	The set of edges in G with both endnodes in X .
$\lambda_G(u, v)$	The minimum number of edge-disjoint paths between nodes u and v .

Directed graphs

$D = (V, A)$	A directed graph (shortly, digraph) on node set V with edge set A .
$\binom{V}{2}$	The set of all (directed) edges on node set V .
$\rho_D(X)/\delta_D(X)$	The number of edges in D entering/leaving node set X .
$\delta_D(X, Y)$	The number of directed edges in D from $X - Y$ to $Y - X$.
$d_D(X, Y)$	$= \delta_D(X, Y) + \delta_D(Y, X)$.
$\overline{d}_D(X, Y)$	$= \delta_D(X \cap Y, V - (X \cup Y)) + \delta_D(V - (X \cup Y), X \cap Y)$.

Set pairs

$K = (K^-, K^+)$	A set pair (see Section 1.1).
$\mathcal{S} = \mathcal{S}_V$	The set of all set pairs on node set V .
$\delta_F(K)$	The number of edges in edge set F covering K .
$K \preceq L$	$K^- \subseteq L^-$ and $K^+ \supseteq L^+$.
$K \wedge L$	$= (K^- \cap L^-, K^+ \cup L^+)$ for dependent set pairs K and L .
$K \vee L$	$= (K^- \cup L^-, K^+ \cap L^+)$ for dependent set pairs K and L .
$\mathcal{O} = \mathcal{O}_D$	The set of one-way pairs in the digraph D .
$\mathcal{O}^1 = \mathcal{O}_D^1$	The set of strict one-way pairs in the $(k - 1)$ -connected digraph D .
$s(K)$	$= V - (K^- \cup K^+) $.

Miscellaneous

$\mathbb{Z}_+/\mathbb{R}_+$	the set of nonnegative integer/real numbers.
x^+	$= \max\{0, x\}$, for a number $x \in \mathbb{R}$.
$f(Z)$	$= \sum_{z \in Z} f(z)$ for a vector $f : V \rightarrow \mathbb{R}$ and a subset $Z \subseteq V$.
$X + v$	$= X \cup \{v\}$ for $X \subseteq V, v \in V$.
$X - v$	$= X - \{v\}$ for $X \subseteq V, v \in V$.
X intersects Y	$X \cap Y, X - Y, Y - X$ are all nonempty for $X, Y \subseteq V$.
X crosses Y	$X \cap Y, X - Y, Y - X, V - (X \cup Y)$ are all nonempty for $X, Y \subseteq V$.
X is an $u\bar{v}$ -set	For $X \subseteq V, u \in X$ and $v \notin X$; used also for more than two nodes.
$x \prec y$	$x \preceq y$ and $x \neq y$ for a partial order \preceq .
$\bigcup \mathcal{X}$	$= \bigcup_{i=1}^t X_i$ for a subpartition $\mathcal{X} = (X_1, \dots, X_t)$.
$p(\mathcal{X})$	$= \sum_{i=1}^t p(X_i)$ for $p : 2^V \rightarrow R$ and a subpartition $\mathcal{X} = (X_1, \dots, X_t)$.

Acknowledgement

It has been a real privilege being a student of András Frank. His profound knowledge of combinatorial optimization is combined with inquisitive curiosity and generous attitude. Besides several intriguing questions that motivated all results of the thesis substantially, I could learn the importance of the deeper understanding of problems from him. Also, his uncompromising aesthetic standards concerning the clarity and simplicity of proofs and presentation have had a great impact on me, even if it is not really apparent from this thesis.

András' great achievement is the foundation of the EGRES group, a unique opportunity for facilitating discussions and for contemplating maths together with excellent colleagues. Let me mention Tamás Király first, his prudence and helpfulness, and his infinite patience in listening to my ideas and carefully reading almost all of my papers. I am also grateful to my co-authors Erika Kovács, Kristóf Bérczi and András Benczúr; it was great working together with them. I have learned a lot from the courses of Tibor Jordán and Zoli Király and from discussions with them. I have been close friends with several members of the EGRES for a long time. With Juli Pap, we have done all our studies together since the age of twelve. Gyuszkó Pap and Marci Makai have always served as examples for me and had a large impact on my decisions on studying at ELTE and joining the EGRES. The helpfulness and kind personality of Jácint Szabó and Attila Bernáth have been contributing a lot to the splendid atmosphere of the group. It has been a great, but unfortunately rare experience doing maths together with Misi Bárász. Tamás Fleiner is a nice and skillful person both in mathematical and real-world problems, being the altruist bicycle repair main of the maths community. Balázs Fleiner is the most scrupulous reader I have ever met, I am grateful to him for the accurate reading of the result of Chapter 3.

I am also grateful to all anonymous referees of my papers, and to Attila, Gyuszkó and Tamás for some remarks and suggestions on this thesis. I would like to thank in advance the efforts of the referees and any other possible readers; I apologize that it grew into something definitely hard to read.

I have traveled a lot during my long-lasting PhD studies. While not always directly connected to my thesis, these stays gave me a superb opportunity of working together with mathematicians from all over the world. First of all, I would like to express my gratitude to Siemens AG and Zuse Institute for providing me a three-year scholarship; their financial support meant a great help. Also, I had the opportunity to spend three months working in both places to get acquainted with applied mathematics. I wish to thank Michael Hofmeister and Christian Royer at Siemens, and Martin Grötschel, Tobias Harks and Ambros Gleixner at ZIB for their assistance and hospitality. I am also grateful to Microsoft Research, in particular, to Laci Lovász for hosting me for one month in Redmond in 2004. I enjoyed a lot the collaborations with Santosh Vempala and Misi Bárász here and next year at Georgia Tech. I am indebted to Martin Loeb for his kindness and for our delightful collaborations during my stays at Charles University in Prague over the

years. Let me also thank the hospitality of András Sebő and Zoli Szigeti in Grenoble, where I spent three months in 2007, supported by the European MCRTN ADONET.¹

From 2004 to 2007, I was funded by the ELTE PhD scholarship. In the last four months of 2007 I was employed as a research assistant at the Operations Research Department of ELTE, funded by the Hungarian National Foundation for Scientific Research Grant (OTKA). Since January 2008, I have been employed by the MTA-ELTE Egerváry Research Group (EGRES), funded by the Hungarian Academy of Sciences. Currently, I am also participant in the Deák Ferenc Scholarship Programme of the Hungarian Ministry of Education. During the last several years, I also received support - besides the aforementioned Siemens-ZIB Scholarship and ADONET Grant - from the OTKA² and from the France Télécom.

I am grateful to Zsóka Kosztolányiné Nagy, my high-school maths teacher, for all her efforts and patience. For several years, I was fortunate to participate in the marvellous mathematical camps of Lajos Pósa, which have been a crucial experience. I am grateful to all my professors at ELTE during my studies.

I lived in Eötvös Collegium for eight years, starting from my first undergraduate year in 1999. I am grateful to all my friends here for the great atmosphere.

I would like to thank my parents, László Végh and Margit Csatlós, and my sisters, Judit and Zsuzsi for all their love and support. Finally, I would like to express my greatest thanks to my love Bori for all the years we have spent together and for her patience during the completion of this thesis.

¹Grant Number 504438

²Grant Numbers K60802, T037547, and TS049788.

Chapter 1

Introduction

The first family of problems considered in the thesis is connectivity augmentation. Given a graph and a positive integer k , we want to find a minimum number of edges whose addition results in a k -node-connected or k -edge-connected digraph. Both edge- and node-connectivity augmentation can be considered in both directed and undirected graphs, which raises four different questions, revealing essential differences both in terms of difficulty and of applicable techniques. An important special case is augmenting connectivity by one, that is, when the input graph is assumed to be already $(k - 1)$ -edge- or node-connected.

A practical motivation is survivable network design. In a network (e.g computer or telecommunication network, electric power supply network), it is utterly important to maintain a path between any two nodes. k -node- or k -edge-connectivity of a graph can be interpreted in terms of security: the network remains connected even if arbitrary $k - 1$ nodes or edges are removed due to attack or failure. In the connectivity augmentation problem, we want to increase the security of an already existing network by adding new connections. From a practical point of view, a minimum cost solution is more desirable: adding different edges may have different costs, and we want to find a minimum cost augmenting edge set. Unfortunately, this problem is NP-complete even in the simplest cases.

Somewhat surprisingly, the cardinality versions turned out to be polynomial time solvable in three of the four basic problems. Undirected edge-connectivity augmentation was solved by Watanabe and Nakamura in 1987 [75], directed edge-connectivity by Frank in 1992 [23], and directed node-connectivity by Frank and Jordán in 1995 [31]. The complexity of undirected node-connectivity augmentation has been a longstanding open question in combinatorial optimization.

For both undirected and directed edge-connectivity augmentation, relatively simple min-max formulae hold. The dual optimum value is given by a partition of the nodes and can be determined via an essentially greedy algorithm. The key technique here is splitting off: Lovász' theorem for undirected and Mader's theorem for directed graphs. In the case of undirected

edge-connectivity, far-reaching generalizations are made possible by Mader's powerful splitting off theorem on preserving local edge-connectivity. Using this theorem, Frank solved local edge-connectivity augmentation, the problem with possibly different connectivity requirements for any pair of nodes. Chapter 5 contains new proofs to classical theorems in this field using the technique of edge-flippings. It also gives partial results towards a generalization, when new edges may only be added between different classes of a fixed partition of the nodes.

For directed node-connectivity augmentation, the dual optimum cannot be described simply by partitions. The novel contribution of Frank and Jordán [31] is the introduction of set pairs. They presented a general abstract theorem (Theorem 1.1) on covering positively crossing supermodular functions on set pairs. The theorem is applicable, among other problems, to directed node-connectivity augmentation. Also, the proof is based on the classical uncrossing technique and it is astonishingly simple. They also gave a polynomial time algorithm for finding an optimal solution. However, their algorithm strongly relied on the ellipsoid method, and thus the question of finding a purely combinatorial algorithm remained open. In Chapter 2 we present such an algorithm, a joint work with András Frank, for augmenting connectivity by one. As one of the main results of the thesis, Chapter 4 provides a completely different type of combinatorial algorithm for the general augmentation problem, a joint result with András Benczúr. It also gives a new, algorithmic proof of Theorem 1.1.

As already mentioned, the complexity status of undirected node-connectivity augmentation is still open. In Chapter 3 we prove a min-max formula for the important special case of augmenting connectivity by one, settling a conjecture of Frank and Jordán from 1994. We also give combinatorial algorithm for finding an optimal solution.

The second main topic of the thesis is constructive characterization, a certain building procedure for describing a class of graphs. A classical example is the ear decomposition of 2-connected graphs. Constructive characterizations are also known for higher connectivity, for example, for k -edge-connected graphs and digraphs. These results are strongly related to the field of connectivity augmentation, with splitting off being the most important method. In Chapter 6, we give a constructive characterization of the so called (k, ℓ) -edge-connected digraphs. This is a joint work with Erika Renáta Kovács and proves a conjecture of András Frank. Our result gives a common generalization of a number of previously known characterizations, and naturally fits into the framework defined by splitting off and orientation theorems.

The rest of this chapter is organized as follows. In Sections 1.1-1.4 we exhibit the background of our results. First, Section 1.1 presents Theorem 1.1 on covering positively crossing supermodular functions along with its main applications. Section 1.2 gives an overview of previous connectivity augmentation algorithms. Section 1.3 and Section 1.4 are devoted to the fields of local edge-connectivity and constructive characterizations, respectively. There is a broad literature on each of these topics and we do not intend to give comprehensive overviews here, but

restrict ourselves to concepts and theorems in direct connection to the results of the thesis.¹ The core of the entire thesis is Section 1.5, where we state the main results of each chapter, sketch the main ideas of the proofs and point out the connections between different chapters.

1.1 The Frank-Jordán Theorem and node-connectivity augmentation

Let us call $K = (K^-, K^+)$ a **set pair** if K^- and K^+ are disjoint nonempty subsets of the ground set V . K^- is called the **tail** and K^+ the **head** of K . Let \mathcal{S} denote the set of all set pairs. We say that a (directed) edge $xy \in V^2$ **covers** the pair K if $x \in K^-$, $y \in K^+$.²

Two set pairs $K = (K^-, K^+)$ and $L = (L^-, L^+)$ are **tail-disjoint** if $K^- \cap L^- = \emptyset$, **head-disjoint** if $K^+ \cap L^+ = \emptyset$, and **independent** if they are either tail- or head-disjoint. This is equivalent to the property that no edge in V^2 covers both K and L . Two non-independent set pairs are called **dependent**. A set \mathcal{F} of set pairs is **independent** if its members are pairwise independent.

A natural partial order on \mathcal{S} can be defined as follows: $K \preceq L$ if $K^- \subseteq L^-$ and $K^+ \supseteq L^+$. The pairs K and L are **comparable** if $K \preceq L$ or $L \preceq K$. Two dependent, but not comparable pairs are called **crossing**.

For dependent K and L , let us define the set pairs $K \wedge L = (K^- \cap L^-, K^+ \cup L^+)$ and $K \vee L = (K^- \cup L^-, K^+ \cap L^+)$. For the partial order \preceq , $K \wedge L$ is the unique greatest common lower bound and $K \vee L$ the least common upper bound. Nevertheless, (\mathcal{S}, \preceq) is not a lattice since $K \vee L$ and $K \wedge L$ are defined only for dependent set pairs.

The non-negative integer valued function p on \mathcal{S} is called **positively crossing supermodular** if

$$p(K) + p(L) \leq p(K \wedge L) + p(K \vee L)$$

whenever $K, L \in \mathcal{S}$, K and L are dependent and $p(K), p(L) > 0$.

For a multiset F consisting of edges in V^2 and a set pair $K \in \mathcal{S}$, let $\delta_F(K)$ denote the number of edges in F covering K . We say that the edge set F **covers** the function p if $\delta_F(K) \geq p(K)$ for every set pair $K \in \mathcal{S}$. Let τ_p denote the minimum size of an edge set covering p , and let $\nu_p = \max\{\sum_{K \in \mathcal{F}} p(K) : \mathcal{F} \text{ independent}\}$. $\nu_p \leq \tau_p$ clearly holds, since an edge may cover at most one member of an independent system. The following theorem states that this in fact holds with equality:

¹We use several well-known results (e.g. Mengers's and Dilworth's theorems, the König-Hall or Berge-Tutte theorems) without references. For all such theorems we refer the reader to Schrijver's monography [69].

²By V^2 we denote the set of all directed edges on a ground set V , while $\binom{V}{2}$ stands for the set of all undirected edges on V .

Theorem 1.1 (Frank and Jordán, 1995 [31]). *Given a ground set V and a positively crossing supermodular function p on the set pairs, $\tau_p = \nu_p$.*

Before turning to the applications, let us consider the important special case when p takes values only 0 and 1. Let $\mathcal{S}_1 = \{K \in \mathcal{S} : p(K) = 1\}$. The supermodularity of p implies that if $K, L \in \mathcal{S}_1$ are dependent then $K \wedge L, K \vee L \in \mathcal{S}_1$. A family of set pairs satisfying this property is called **crossing**. In fact, we may obtain every crossing family in this form. Given a crossing family \mathcal{F} , the function p defined by $p(K) = 1$ if $K \in \mathcal{F}$ and $p(K) = 0$ if $K \notin \mathcal{F}$ is positively crossing supermodular. This observation leads to the following corollary of Theorem 1.1. For a crossing family \mathcal{F} , let $\tau(\mathcal{F})$ denote the minimum number of edges covering \mathcal{F} , and let $\nu(\mathcal{F})$ be the maximum number of pairwise independent members of \mathcal{F} .

Theorem 1.2. *Given a crossing family \mathcal{F} of set pairs, $\nu(\mathcal{F}) = \tau(\mathcal{F})$.*

Let us now exhibit some applications of Theorem 1.1, starting with the most prominent one, directed connectivity augmentation.

1.1.1 Directed connectivity augmentation

We commence by giving the precise definition of k -edge- and node-connectivity. All directed and undirected graphs in the thesis will be allowed to have parallel edges and loops. By edge set we will always mean a multiset of edges, even if not mentioned explicitly. A directed graph is called **strongly connected** if it contains a directed path between any two nodes. An undirected or directed graph is called **k -node-connected** or shortly, **k -connected** if the number of nodes is at least $k + 1$, and after the deletion of any subset of at most $k - 1$ nodes, the remaining graph is still connected if undirected, and strongly connected if directed. Analogously, an undirected or directed graph is called **k -edge-connected**, if after the deletion of any at most $k - 1$ edges, the remaining graph is still (strongly) connected. It is well-known, by versions of Menger's theorem, that a graph or digraph is k -node-connected (respectively, k -edge-connected) if and only if there are k internally node-disjoint (edge-disjoint) paths from each node to every other node (and the graph has at least $k + 1$ nodes in the k -node-connected case).

In the directed node-connectivity augmentation problem we are given a digraph $D = (V, A)$ and a target value k , and we want to add a minimum number of new edges to D to make it k -connected. A set pair $K \in \mathcal{S}$ is called a **one-way pair** if $\delta_D(K) = 0$, that is, there are no edges in D covering K . We denote by $\mathcal{O} = \mathcal{O}_D$ the set of one-way pairs. For a set pair K , let us define $s(K) := |V - (K^- \cup K^+)|$. The following simple claim shows that we may restrict our attention to the one-way pairs:

Claim 1.3 ([31]). *D is k -connected if and only if $s(K) \geq k$ for every $K \in \mathcal{O}$.*

Let us define the function p as follows: $p(K) := (k - s(K))^+$ if $K \in \mathcal{O}$, and $p(K) := 0$ if $K \notin \mathcal{O}$. It is easy to verify that p is positively crossing supermodular. By the previous claim, $D + F$ is k -connected if and only if F covers p . Hence Theorem 1.1 specializes to:

Theorem 1.4. *For a digraph $D = (V, A)$, the minimum number of edges whose addition makes D k -connected equals the maximum value of $\sum_{i=1}^{\ell} (k - s(K_i))$ over pairwise independent one-way pairs K_1, \dots, K_{ℓ} .*

Assume now that the digraph D is already $(k - 1)$ -connected, implying $s(K) \geq k - 1$ for all one-way pairs. We call a one-way pair **strict** if $s(K) = k - 1$ and denote their set by $\mathcal{O}^1 = \mathcal{O}_D^1$. The theorem simplifies to the following form:

Theorem 1.5. *For a $(k - 1)$ -connected digraph $D = (V, A)$, the minimum number of edges whose addition makes D k -connected equals the maximum number of pairwise independent strict one-way pairs.*

In Chapter 2, we will also use the following mild generalization of Theorem 1.5. This is also a simple consequence of Theorem 1.2.

Theorem 1.6. *For a $(k - 1)$ -connected digraph $D = (V, A)$, let $\mathcal{F} \subseteq \mathcal{O}_D^1$ be a crossing family of strict one-way pairs. Then $\nu(\mathcal{F}) = \tau(\mathcal{F})$.*

1.1.2 Other applications

Győri's theorem

Perhaps the most astonishing applications of Theorem 1.1 are Győri's theorems on generators of interval systems and on rectangle coverings. Let us start with the first problem: let \mathcal{I} be a finite set of closed intervals in $[0, 1]$. We say that the set \mathcal{B} of closed intervals **generates** \mathcal{I} if every interval in \mathcal{I} is the union of some members of \mathcal{B} . (For example, \mathcal{I} generates itself.) Given \mathcal{I} , we are interested in the minimum size of a set generating it. For an $I \in \mathcal{I}$ and an interior point $x \in I$, we say that (I, x) is a represented interval. Two represented intervals (I, x) and (J, y) are called **independent** if $I \cap J$ does not contain both x and y .

Theorem 1.7 (Győri, 1984 [38]). *The minimum size of a generator of a set \mathcal{I} equals the maximum number of pairwise independent represented intervals in \mathcal{I} .*

This was originally conjectured by Frank in the late seventies and proved by Győri in 1984. Győri's original proof was quite sophisticated and the theorem did not show any relations to other min-max theorems known by that time. Let us now derive this result from Theorem 1.2. It is clear that $[0, 1]$ can be replaced by a path $P = \{v_1, e_1, v_2, e_2, \dots, e_{t-1}, v_t\}$ with nodes v_i and edges e_i . The intervals correspond to subpaths of \mathcal{P} . For a path $I = \{v_h, e_h, \dots, e_{k-1}, v_k\}$ and

an edge e_i with $h \leq i \leq k - 1$, we may define the set pair $K_{I,e_i} = (\{v_h, \dots, v_i\}, \{v_{i+1}, \dots, v_k\})$. Finding a system of generators is equivalent to covering the set pairs K_{I,e_i} for every possible choice of I and e_i . It is easy to verify that these set pairs form a crossing system, and two pairs K_{I,e_i} and K_{J,e_j} are independent if and only if (I, x) and (J, y) are independent for any interior points $x \in e_i, y \in e_j$. Theorem 1.1 also easily implies an extension of Theorem 1.7 for intervals on a circuit instead of intervals in $[0, 1]$; this generalization could not be obtained from Györi's original proof.

The theorem has a nice application in combinatorial geometry. We say that a polygon in the plane is **rectilinear** if all edges are vertical and horizontal lines. A rectilinear polygon is **vertically convex** if its intersection with every vertical line is an interval. For a rectilinear polygon R , we say that \mathcal{H} is a **rectangle cover** of R if \mathcal{H} is a set of rectangles contained in R whose union is R . A set P of points in R is called **independent** if no two points in P can be covered by a rectangle contained in R .

Theorem 1.8. *For a vertically convex rectilinear polygon R , the minimum size of a rectangle cover of R equals the maximum size of an independent point set in R .*

K_{tt} -free t -factors in bipartite graphs

Given an undirected graph $G = (V, E)$, a natural relaxation of the Hamiltonian cycle problem is to find a $C_{\leq k}$ -free 2-matching, that is, a subgraph with maximum degree 2 containing no cycle of length at most k . Cornuéjols and Pulleyblank [14] showed this problem to be NP-complete for $k \geq 5$. In his Ph.D. thesis [40], Hartvigsen proposed a solution for the case $k = 3$. The case $k = 4$ is still open along with the other natural question of finding a maximum C_4 -free 2-matching (possibly containing triangles). Only some partial results are known so far (see [7] and [8]).

However, the C_4 -free 2-matching problem turns out to be tractable under the assumption that G is bipartite. This was solved by Hartvigsen [41, 42] and Király [53]. A generalization of the problem to maximum $K_{t,t}$ -free t -matchings was given by Frank [27], who observed that this can indeed be deduced from Theorem 1.1.

Theorem 1.9 (Frank, 2003 [27]). *The maximum size of a $K_{t,t}$ -free t -matching of a bipartite graph $G = (S, T; E)$ equals*

$$\min_{Z \subseteq S \cup T} (t|Z| + i(V - Z) - c_t(Z)), \quad (1.1)$$

where $c_t(Z)$ denotes the number of connected components of $(S \cup T) - Z$ which are $K_{t,t}$'s.

Let us define a function p on set pairs on $V = S \cup T$ as follows. If $K^- \subseteq S, K^+ \subseteq T$, and G spans a complete bipartite graph between K^- and K^+ , then let $p(K) = (|K^-| + |K^+| - 2t + 1)^+$

if $|K^-|, |K^+| \geq 2$, and $p(K) = (|K^-| + |K^+| - t - 1)^+$ if $|K^-| = 1$ or $|K^+| = 1$. Let $p(K) = 0$ in all other cases. It can be verified that this function is positively crossing supermodular, and if F is an edge set covering p then $E - F$ is a $K_{t,t}$ -free 2-matching. Moreover, a dual optimal solution may be transformed to the form (1.1).

A generalization of this problem is if we do not exclude all $K_{t,t}$ subgraphs, but only a certain subset of them is forbidden. The above reduction method fails to work, still, Makai [65] generalized Theorem 1.9 for this setting. To this end, he formulated and proved a nontrivial generalization of Theorem 1.1 - which is indeed the only nontrivial generalization known so far. However, this theorem and the other extensions of Theorem 1.9 are beyond the scope of this thesis.

There is an interesting connection between the matching problems above and undirected connectivity augmentation. It is easy to see that for $k = n - 2$ ($n = |V|$), connectivity augmentation is equivalent to finding a maximum matching in the complement graph of G . For $k = n - 3$, the problem is equivalent to finding a maximum C_4 -free 2-matching. However, for $k < n - 3$ the problem corresponding to connectivity augmentation is not $K_{t,t}$ -free t -matchings, but t -matchings not containing any complete bipartite graph $K_{a,b}$ with $a + b = t + 2$. This latter problem can also be solved in bipartite graphs using Theorem 1.1.

k -elementary bipartite graphs

Let $G = (S, T; E)$ be a bipartite graph. It is well known by Hall's theorem that there exists a matching covering S if and only if $|X| \leq |\Gamma(X)|$ holds for every $X \subseteq S$, where $\Gamma(X) \subseteq T$ denotes the set of neighbours of X . G is called **elementary bipartite** if either $|S| = |T| = 1$ and E consists of a single edge or $|S| = |T| > 1$ and the stronger property $|X| + 1 \leq |\Gamma(X)|$ holds for every $\emptyset \neq X \subsetneq S$. This is a well-studied class of graphs, see e.g. [61, Chapter 4].

As a generalization, for $k \in \mathbb{Z}_+$ we say that the bipartite graph $G = (S, T; E)$ is **k -elementary** (with respect to S) if $|X| + k \leq |\Gamma(X)|$ or $\Gamma(X) = T$ for every $\emptyset \neq X \subseteq S$. (Note that $|S| = |T|$ is not being assumed.) The following problem is an analogue of connectivity augmentation. Given a bipartite graph $G = (S, T; E)$, add a minimum number of edges between S and T to get a k -elementary bipartite graph. We say that the set X is **legal** if $\emptyset \neq X \subseteq S$, $\Gamma(X) \neq T$. Two legal sets X and Y are **independent** if either $X \cap Y = \emptyset$ or $\Gamma(X \cup Y) = T$.

Theorem 1.10. *For a bipartite graph $G = (S, T; E)$, the minimum number of edges between S and T whose addition makes G elementary bipartite equals the maximum value of $\sum_{i=1}^t (k + |X_i| - \Gamma(X_i))$ over pairwise independent legal sets X_1, \dots, X_t .*

This can easily be derived from Theorem 1.1 by mapping each legal set X to the set pair $K_X = (X, T - \Gamma(X))$ with $p(K_X) = (k + |X| - \Gamma(X))^+$ and $p(K) = 0$ for any other set pair

K . Clearly, this function is positive crossing supermodular, and the set pairs K_X and K_Y are independent if and only if the legal sets X and Y are independent.

Connectivity augmentation may be easily reduced to this problem. Given the digraph $D = (V, A)$ with $|V| \geq k + 1$, construct a bipartite graph $G = (S, T; E)$ by associating two nodes $v' \in S$ and $v'' \in T$ and an edge $v'v'' \in E$ with each $v \in V$, and furthermore an edge $u'v'' \in E$ with each edge $uv \in A$. This graph is k -elementary bipartite if and only if D is k -connected. A similar reduction is possible in the other direction as well, assuming that $|S| = |T|$ and that G is 0-elementary (that is, it satisfies the Hall-condition). This correspondence will be useful for the algorithmic aspects of augmenting directed connectivity by one in Chapter 2 and even for undirected connectivity augmentation in Chapter 3.

Directed edge-connectivity augmentation

Augmenting directed edge-connectivity is considerably easier than node-connectivity, and was solved in 1992 by Frank [23] via Mader's directed splitting off theorem (Theorem 1.28). In Section 1.3 we show that an analogous argument works out for undirected edge-connectivity augmentation as well.

Let us now formulate the min-max formula and show how it can also be derived from Theorem 1.1.

Theorem 1.11 (Frank, 1992 [23]). *Given a digraph $D = (V, A)$, the minimum number of edges whose addition makes D k -edge-connected equals the maximum value of*

$$\max\left\{\sum_{i=1}^{\ell}(k - \rho(X_i)), \sum_{i=1}^{\ell}(k - \delta(X_i))\right\},$$

over subpartitions $\{X_1, \dots, X_\ell\}$.

Define a positively crossing supermodular function p on \mathcal{S} by giving nonzero values only to set pairs corresponding to cuts, namely, let $p(K) = (k - \rho(K^+))^+$ whenever $K^- \cup K^+ = V$ and $p(K) = 0$ otherwise. Covering p is clearly equivalent to k -edge-connectivity augmentation. The theorem follows by showing that the complex structure of pairwise independent set pairs breaks down to the simple dual optimum in Theorem 1.11, established by the next claim.

Claim 1.12. *If any two among the sets $X_1, \dots, X_\ell \subseteq V$ are disjoint or co-disjoint, then either they are all pairwise disjoint or all pairwise co-disjoint. (Two sets are called co-disjoint if their union is V).* □

In Section 6.3 we present Theorem 6.19, a generalization of this theorem for positively crossing supermodular set functions, derivable from Theorem 1.1 (more precisely, from its degree-prescribed version, which we do not discuss here).

***ST*-edge-connectivity augmentation**

Whereas Theorem 1.11 can also be obtained by the significantly simpler splitting off technique, this does not hold for the following generalization of edge-connectivity augmentation. Let $D = (V, A)$ be a digraph with two (not necessarily disjoint) sets $S, T \subseteq V$. D is called ***k*-*ST*-edge-connected** if for any $s \in S$ and $t \in T - s$, there are at least k -edge-disjoint paths from s to t . $S = T = V$ gives k -edge-connectivity, while $S = \{r_0\}$, $T = V - \{r_0\}$ gives rooted k -edge-connectivity.

The problem of adding a minimum number of edges to D to make it k -*ST*-edge-connected is NP-complete already for $k = 1$. However, if adding new edges only between S and T is allowed, the problem becomes polynomially solvable. Define p on \mathcal{S} to be positive only on set pairs K with $K^- \subseteq S$, $K^+ \subseteq T$. On such pairs, let $p(K) = \max\{(k - \rho(X))^+ : X \cap T = K^+, S - X = K^-\}$. This is a positively crossing supermodular function, and its coverings coincide with the augmenting edge sets consisting of edges from S to T .

We may also give a min-max formula in terms of sets instead of set pairs. Let X be called an *ST*-set if $X \cap T \neq \emptyset$, $S - X \neq \emptyset$. Two *ST*-sets X and Y are called **independent** if either $X \cap Y \cap T = \emptyset$ or $S \subseteq X \cup Y$.

Theorem 1.13. *For a digraph $D = (V, A)$ with $S, T \subseteq V$, the minimum number of edges from S to T whose addition makes D k -*ST*-edge-connected equals the maximum of $\sum_{i=1}^{\ell} (k - \rho(X_i))^+$ over pairwise *ST*-independent *ST*-sets X_1, \dots, X_{ℓ} .*

The reason why this problem is more complicated than edge-connectivity augmentation is that the structure of *ST*-independence cannot be simplified to partitions and co-partitions as in Claim 1.12.

1.2 Previous algorithmic results on connectivity augmentation

For $k = 1$, the notions of 1-edge- and 1-node-connectivity coincide, both giving connectedness in the undirected and strongly connectedness in the directed case. Augmenting an undirected graph to be connected is trivial (and even the minimum cost version is tractable via Kruskal's algorithm). The case $k = 1$ for directed graphs was solved in 1976 by Eswaran and Tarjan [19].

As already mentioned, min-max formulae and polynomial time algorithms for optimal edge-connectivity augmentation were developed by Watanabe and Nakamura in 1987 [75] for the undirected and by Frank in 1992 [23] for the directed case; undirected edge-connectivity will be discussed in Section 1.3.

Concerning directed node-connectivity, even the case $k = 2$ has not been settled until the result of Frank and Jordán in 1995 [31]. The algorithm in their paper strongly relied on the

ellipsoid method, thus finding a combinatorial algorithm remained an open problem. The first result towards this direction was given by Enni in 1999 [18], by nontrivially extending the algorithm of Eswaran and Tarjan for 1-*ST*-edge-connectivity augmentation. For fixed k , Frank and Jordán themselves gave a combinatorial algorithm in 1999 [32] for directed connectivity augmentation - that is, the running time is the product of a polynomial of n and an exponential function of k .

For the 0-1 valued case (Theorem 1.2), two completely different and independent algorithms were given in 2003 by Frank [26] and Benczúr [4]. However, Frank's algorithm was not directly applicable for graph connectivity augmentation. Our joint result with Frank presented in Chapter 2 is an extension of this work. In contrast, the result of Chapter 4 is the extension of the algorithm of Benczúr.

As shown in the previous section, Györi's theorem (Theorem 1.7) is also a special case of Theorem 1.1. Various polynomial time algorithms were given by Franzblau and Kleitman in 1986 [37], by Lubiw in 1990 [62], by Knuth in 1996 [55], by Frank in 1999 [25] and by Benczúr, Király and Förster in 1999 [5]. Some fundamental ideas of [26] (and thus of Chapter 2) derive from [25].

For undirected connectivity augmentation, the situation is radically different. The complexity of the general problem is still unknown; even augmenting by one has been open for a long time. This problem is settled in Chapter 3 of this thesis. In the same paper [19], Eswaran and Tarjan also gave an algorithm for augmenting a graph to be 2-connected. Watanabe and Nakamura solved the case $k = 3$ in 1993 [76] while $k = 4$ was done by Hsu in 2000 [44]. Other solved special cases include $k = n - 2, n - 3$: As mentioned in Section 1.1.2, connectivity augmentation for $k = n - 2$ for the graph G is equivalent to finding a maximum matching in the complement graph of G . Similarly, augmentation by one for $k = n - 3$ is equivalent to finding a maximum square-free 2-matching in a subcubic graph, solved recently by Bérczi and Kobayashi [7].

The best previously known result is due to Jackson and Jordán from 2005 [47]. They gave a polynomial time algorithm for finding an optimal augmentation for any fixed k . The running time is bounded by $O(n^5 + f(k)n^3)$, where $f(k)$ is an exponential function of k . They proved even stronger results for some special classes of graphs: for example, the running time of the algorithm is a polynomial of n if the minimum degree is at least $2k - 2$. An analogous result is by Liberman and Nutov [59]. They gave a polynomial time algorithm for increasing connectivity by one under the assumption that there exists a set $Z \subseteq V$ with $|Z| = k - 1$ so that $G - Z$ has at least k connected components. (It can be decided in polynomial time whether a graph contains such a set, see Cheriyan and Thurimella [13].)

It is straightforward to give a 2-approximation for connectivity augmentation by replacing each edge by two oppositely directed edges and using that directed node-connectivity can be augmented optimally. For augmenting connectivity by one, Jordán [49, 50] gave an algorithm

finding an augmenting edge set larger than the optimum by at most $\lceil \frac{k-2}{2} \rceil$. Jackson and Jordán [46] extended this result for general connectivity augmentation with an additive term of $\lceil \frac{k(k-1)+4}{2} \rceil$. A slightly weaker, similar result was established also by Ishii and Nagamochi [45]. (The running times of these algorithms can be bounded by polynomials of n .)

1.3 Undirected edge-connectivity augmentation

The min-max formula on undirected edge-connectivity augmentation is the following.

Theorem 1.14 (Watanabe and Nakamura, 1987 [75]). *For a graph $G = (V, E)$ and a connectivity requirement $k \geq 2$, the minimum number of edges whose addition makes G k -edge-connected equals the maximum of $\lceil \frac{1}{2} \sum_{i=1}^{\ell} (k - d(X_i)) \rceil$ over subpartitions X_1, \dots, X_ℓ of V .*

In contrast with the other basic augmentation problems, here we can also cope with **local edge-connectivity augmentation**, that is, we may have a different connectivity requirement for each pair of nodes: $r(u, v) = r(v, u)$ for the nodes $u, v \in V$. **Global edge-connectivity augmentation** will refer to the case $r \equiv k$ for some $k \in \mathbb{Z}_+$.

For an undirected graph $G = (V, E)$, let $\lambda(u, v) = \lambda_G(u, v)$ denote the maximum number of edge-disjoint paths between u and v . By Menger's theorem, it is well-known that $\lambda_G(u, v) = \min\{d_G(X) : X \subseteq V, u \in X, v \notin X\}$. Given a function $r : V \times V \rightarrow \mathbb{Z}_+$, we say that $G = (V, E)$ is **r -edge-connected** if $\lambda(u, v) \geq r(u, v)$ for any $u, v \in V$.

F is called an **augmenting edge set** (for G with respect to r) if $G + F$ is r -edge-connected. This can be equivalently formulated in terms of cuts: let $R(\emptyset) = R(V) = 0$,

$$R(X) := \max\{r(u, v) : u \in X, v \notin X\} \text{ if } \emptyset \neq X \subsetneq V, \quad (1.2)$$

and let $p(X) := (R(X) - d_G(X))^+$. Then $G + F$ is r -edge-connected if and only if

$$d_F(X) \geq p(X) \text{ for every } X \subseteq V. \quad (1.3)$$

For an arbitrary set function p , we say that the edge set F **covers** p if (1.3) holds. Frank's following theorem gives a min-max formula on the minimum size of an augmenting edge set. For a partition $\mathcal{X} = \{X_1, \dots, X_\ell\}$, let $p(\mathcal{X}) = \sum_{i=1}^{\ell} p(X_i)$. A set $C \subseteq V$ is called a **marginal set**, if $R(C) \leq 1$ and $d(C) = 0$.

Theorem 1.15 (Frank, 1992 [23]). *Assume we are given a graph $G = (V, E)$ and the requirement function r so that G contains no marginal sets. Then the minimum number of edges whose addition makes G r -edge-connected equals the maximum value of $\lceil \frac{1}{2} p(\mathcal{X}) \rceil$ over subpartitions \mathcal{X} of V .*

The $\max \leq \min$ direction is clear since we need to add at least $p(X_i) = (R(X_i) - d(X_i))^+$ new edges for each class X_i of \mathcal{X} and a new edge may cover at most two X_i 's. Actually, Frank's original theorem is slightly stronger by excluding only marginal components instead of marginal sets. A connected component $C \subseteq V$ is called a **marginal component** if $R(C) \leq 1$, and $p(U) = 0$ for any $U \subsetneq C$. However, this original version can be easily derived from Theorem 1.15. Also, all subsequent theorems where marginal sets are excluded can be strengthened to exclude only marginal components; we stick to marginal sets for the sake of minor simplifications in some proofs. The condition excluding marginal sets or components is necessary since a graph $G = (V, \emptyset)$ with $r \equiv 1$ needs at least $|V| - 1$ new edges, although the dual optimum is $\left\lfloor \frac{|V|}{2} \right\rfloor$. Nevertheless, even the most general case without any restriction on r can be deduced from Frank's original theorem (and thus from Theorem 1.15), see in [23].

The nontrivial direction is proved via Mader's splitting off theorem, an extremely powerful tool for edge-connectivity problems. By **splitting off** edges $e = xz$ and $f = zy$ we mean the operation of deleting e and f and adding the new edge xy (literally the same definition is used for digraphs as well, see in Section 1.4). We say that a splitting off is **admissible** if for any two nodes $u, v \in V - z$, the local edge-connectivity value $\lambda(u, v)$ does not decrease. The pair of edges xz, zy is **splittable** if splitting off xz and zy is admissible.

Theorem 1.16 (Mader, 1978 [63]). *Let $G = (V + z, E)$ be a graph with $d(z) \neq 3$ so that there is no cut edge incident to z . Then there exist a splittable pair of edges incident to z .*

Based on this theorem, Theorem 1.15 can be deduced via the following intermediate theorem. A $V \rightarrow \mathbb{Z}_+$ function m is called a **degree-prescription** if $m(V)$ even. For a degree-prescription m , an edge set F is called **m -prescribed** if $d_F(v) = m(v)$ for every $v \in V$. Clearly, such an edge set always exists.

Theorem 1.17 ([23]). *Assume we are given a graph $G = (V, E)$ containing no marginal sets, a requirement function r and a degree-prescription m . Then there exists an m -prescribed edge set F so that $G + F$ is r -edge-connected if and only if*

$$m(X) \geq p(X) \quad \forall X \subseteq V. \tag{1.4}$$

This can be proved by adding a new node z to the graph G , and connecting it to each node v by $m(v)$ parallel edges. The resulting graph is r -edge-connected in V and has no cut edges incident to z , hence the iterative application of the splitting off theorem yields the desired F .

By **parity adjusting** of a function $m : V \rightarrow \mathbb{Z}_+$ we mean the following operation: if $m(V)$ is odd then we increase $m(v)$ by one for an arbitrary $v \in V$. The following can be proved using the uncrossing technique (see the detailed argument in Section 5.1.1). If we take an arbitrary m which is a minimal one satisfying (1.4), and furthermore we apply parity adjusting on m , then $m(V)$ will be twice the maximum value in Theorem 1.15. The key property of R we use both in the proof of Theorem 1.16 and in the uncrossing method is that it is **skew supermodular**:

Claim 1.18 ([66],[23]). *For any two subsets $X, Y \subseteq V$, at least one of the following two inequalities hold:*

$$R(X) + R(Y) \leq R(X \cup Y) + R(X \cap Y) \quad (1.5a)$$

$$R(X) + R(Y) \leq R(X - Y) + R(Y - X) \quad (1.5b)$$

This easily implies that the function p is also positively skew supermodular, that is, at least one of the two inequalities hold for p in place of R for any sets X, Y with $p(X), p(Y) > 0$. For the function R , an even stronger property can also be easily verified:

Claim 1.19. *If one of (1.5a) and (1.5b) does not hold, then the other is true with equality.*

For global edge-connectivity augmentation, Theorem 1.16 was preceded by Lovász' global splitting off theorem preserving k -edge-connectivity [60], and Theorem 1.14 was proved based on this theorem. The splitting off technique is also important in context of directed edge-connectivity, discussed in Section 1.4.

Positively crossing supermodular functions

One might wonder if Theorem 1.15 extends to a general covering theorem for arbitrary functions p satisfying certain properties. Unfortunately, the symmetry and positively skew-supermodularity are not enough by themselves: a special case of this problem, local edge-connectivity augmentation of hypergraphs is NP-complete, see [54].

An abstract extension of Theorem 1.14 on global edge-connectivity augmentation was formulated by Benczúr and Frank in 1999 [6], by replacing $k - d(X)$ with a certain type of function $p(X)$. Let $p : 2^V \rightarrow \mathbb{Z}_+$ be an arbitrary symmetric and positively crossing supermodular function, that is, $p(X) = p(V - X)$ for any $X \subseteq V$ and

$$p(X) + p(Y) \leq p(X \cup Y) + p(X \cap Y)$$

holds whenever $p(X), p(Y) > 0$ and X and Y are crossing ($X \cap Y, X - Y$ and $Y - X$ are all nonempty sets and $X \cup Y \neq V$). Note that this also implies

$$p(X) + p(Y) \leq p(X - Y) + p(Y - X)$$

if $p(X), p(Y) > 0$. Theorem 1.14 does not remain true by simply replacing $k - d(X)$ by $p(X)$ and using the subpartition bound $\max \lceil \frac{1}{2} p(\mathcal{X}) \rceil$. In fact, a new type of obstacle should also be taken into account. Let us call a partition $\mathcal{P} = \{X_1, \dots, X_t\}$ of the node set V **p -full** if $p(\bigcup_{i \in I} X_i) > 0$ holds for any nonempty subset $I \subsetneq \{1, 2, \dots, t\}$. Clearly, at least $t - 1$ edges are needed to cover such a p . The maximum cardinality of a p -full partition is called the **dimension of p** and is denoted by $\dim(p)$. While the definition contains exponentially many conditions, the following simple lemma shows that p -fullness can be verified effectively:

Lemma 1.20 ([6]). *Assume that for $\mathcal{P} = \{X_1, \dots, X_t\}$, $p(X_1) = 1$ and $p(X_1 \cup X_i) \geq 1$ for any $i > 0$. Then \mathcal{P} is p -full.*

The theorem is as follows:

Theorem 1.21 (Benczúr and Frank, 1999 [6]). *Let $p : 2^V \rightarrow \mathbb{Z}_+$ be a symmetric positively crossing supermodular function. Then the minimum cardinality of an edge set F covering p is equal to*

$$\max\{\dim(p) - 1, \max \left\lceil \frac{1}{2}p(\mathcal{X}) \right\rceil\},$$

where the second maximum ranges over subpartitions \mathcal{X} of V .

An important application of this theorem is global edge-connectivity augmentation of hypergraphs, solved by Bang-Jensen and Jackson in 1999 [3]. Recall that Theorem 1.15 on local edge-connectivity augmentation was a consequence of the degree-prescribed Theorem 1.17. Similarly, Theorem 1.21 is an easy consequence of the degree-prescribed version.

Theorem 1.22 ([6]). *Let us be given a symmetric positively crossing supermodular function $p : 2^V \rightarrow \mathbb{Z}_+$ and a degree-prescription m . There exists an m -prescribed edge set F covering p if and only if (1.4) holds and furthermore*

$$m(V) \geq \dim(p) - 1. \tag{1.6}$$

A directed counterpart of this theorem is Theorem 6.19. The symmetry of p is not required in that case, and also no obstacle similar to p -full partitions occur.

Partition-constrained problems

The central problem investigated in Chapter 5 is **partition-constrained local edge-connectivity augmentation (PCLECA)**. Given a partition $\mathcal{Q} = (Q_1, \dots, Q_t)$ of V , an edge is called **\mathcal{Q} -legal** if its endnodes lie in different classes of \mathcal{Q} . Given a requirement function r and a partition \mathcal{Q} , we want to find a minimum cardinality set F consisting of \mathcal{Q} -legal edges so that $G + F$ is r -edge-connected.

For global edge-connectivity ($r \equiv k \geq 2$) this problem was solved by Bang-Jensen, Gabow, Jordán and Szigeti [2]. Given a graph $G = (V, E)$, a partition \mathcal{Q} of the nodes and a connectivity requirement $k \geq 2$, let $OPT_{\mathcal{Q}}^k$ denote the minimum number of \mathcal{Q} -legal edges whose addition makes G k -edge-connected. Clearly, the problem is equivalent to covering the function $p(X) = (k - d(X))^+$ by a minimum number of \mathcal{Q} -legal edges.

A natural lower bound on this is the one in Theorem 1.14, namely, $\alpha(G) = \max \left\lceil \frac{1}{2}p(\mathcal{X}) \right\rceil$ over subpartitions \mathcal{X} of V . For a similar bound for each $1 \leq j \leq t$, let us call \mathcal{X} a j -subpartition, if \mathcal{X} is a subpartition of Q_j . Let $\beta_j(G) = \max p(\mathcal{X})$ over j -subpartitions \mathcal{X} . Let $\Psi_{\mathcal{Q}}(G)$ denote the maximum of $\alpha(G)$ and $\beta_j(G)$ for $j = 1, \dots, t$. The theorem is the following.

Theorem 1.23 ([2]). *Given an undirected graph $G = (V, E)$, a partition \mathcal{Q} of the nodes and a connectivity requirement $k \geq 2$, $OPT_{\mathcal{Q}}^k = \Psi_{\mathcal{Q}}(G)$ if k is even, or k is odd and G contains neither a C_4 nor a C_6 -configuration. Otherwise, $OPT_{\mathcal{Q}}^k = \Psi_{\mathcal{Q}}(G) + 1$.*

We define only C_4 -configurations here as we will not need C_6 -configurations in the sequel. For subpartitions \mathcal{Z} and \mathcal{W} , we say that \mathcal{Z} is a **refinement** of \mathcal{W} if each class of \mathcal{Z} is a subset of some class of \mathcal{W} .

Let $\{A_1, A_2, C_1, C_2\}$ be a partition of V , and for some $1 \leq h \leq t$, let \mathcal{Z} be a h -partition which is a refinement of $\{C_1, C_2\}$. These form a C_4 -**configuration** if they fulfil the following: (i) $p(\mathcal{Z}) = \Psi_{\mathcal{Q}}(G)$; (ii) $d_G(C_1, C_2) = d_G(A_1, A_2) = 0$, and (iii) $p(C_j) = \sum\{p(Z) : Z \in \mathcal{Z}, Z \subseteq C_j\}$ for $j = 1, 2$.

Let us see an example: consider $G = (V, E)$ on the node set $V = \{a_1, c_1, a_2, c_2\}$ and edge set $E = \{a_1c_1, c_1a_2, a_2c_2, c_2a_1\}$ (a square). Let $\mathcal{Q} = (\{a_1, a_2\}, \{c_1, c_2\})$ and $k = 3$. At least three new \mathcal{Q} -legal-edges are needed for the augmentation, while $\Psi_{\mathcal{Q}}(G) = 2$.

Similarly to the previous theorems, this one was also proved using splitting off techniques, and a degree-prescribed variation can also be formulated. The proof starts by adding a new node z and an edge set H incident to z with $|H| = \Psi_{\mathcal{Q}}(G)$. (By choosing this edge set, the partition \mathcal{Q} should also be taken into account). A pair of edges xz and yz is called **\mathcal{Q} -legal** if x and y lie in different classes of \mathcal{Q} . As long as possible, we split off \mathcal{Q} -legal admissible pairs of edges incident to z . If all edges incident to z can be removed in such pairs then we have found an optimal \mathcal{Q} -legal augmentation. If not, then either we are able to achieve a complete splitting after undoing one of the previously performed splitting off operations, or the existence of a C_4 - or C_6 -configuration can be verified.

In Chapter 5, we give new proofs of Theorems 1.17 and 1.21 using edge-flippings instead of splitting off. Furthermore, partial results are presented towards the generalization of Theorem 1.23 to local edge-connectivity augmentation. A common generalization of Theorems 1.21 and 1.23 was given by Bernáth, Grappe and Szigeti [11]. A detailed discussion of these topics among plenty of new extensions can be found in the recent thesis of Bernáth [9].

1.4 Constructive characterizations

By a constructive characterization of a graph property \mathcal{P} we mean a set of operations preserving property \mathcal{P} , so that each graph with property \mathcal{P} can be obtained by a sequence of such operations starting from a small set of basic instances. Such characterizations are often useful for proving further properties of graphs with property \mathcal{P} . The following ear decompositions of 2-connected and 2-edge-connected graphs are among the first examples of constructive characterizations.

Proposition 1.24. (i) [77] *An undirected graph is 2-connected if and only if it can be built up from a circuit by iteratively adding new paths whose endpoints are distinct old nodes.*

(ii) [60, Problem 6.28] *An undirected graph is 2-edge-connected if and only if it can be built up from a single node by iteratively adding new paths whose endpoints are (possibly coincident) existing nodes.*

In this section, we focus on results related to higher edge-connectivity. Although the ear decompositions above are almost identical for node- and edge-connectivity, very little is known on characterizing k -node-connected graphs: there are different constructive characterizations for $k = 3$, but none for $k \geq 4$. A survey on constructive characterizations in combinatorial optimization can be found in [57].

An immediate application of Proposition 1.24(ii) is the following. Given an undirected graph G , we want to find a strongly connected orientation of G . A trivial necessary condition is that G should be 2-edge-connected. Using the characterization, sufficiency is also straightforward: when adding a path, let us orient all its edges in the same direction. We will see orientation results for higher edge-connectivity as well and their relation to constructive characterizations. For $2k$ -edge-connected graphs, Lovász proved the following.

Theorem 1.25 (Lovász, 1976 [60, Problem 6.52]). *An undirected graph is $2k$ -edge-connected if and only if it can be obtained from a single node by iteratively applying the following two operations:*

- (i) *add a new edge (possibly a loop),*
- (ii) *subdivide k existing edges and identify the subdividing nodes.*

It is easy to see the equivalence between the case $k = 1$ and the ear decomposition in Proposition 1.24(ii). Mader gave a similar characterization for $2k + 1$ -edge-connected graphs [63]. As for the $k = 1$ case, Theorem 1.25 immediately implies the weak version of Nash-Williams' orientation theorem:

Theorem 1.26 (Nash-Williams, 1960 [66]). *An undirected graph has a k -edge-connected orientation if and only if it is $2k$ -edge-connected.*

A directed counterpart of Theorem 1.25 is due to Mader:

Theorem 1.27 (Mader, 1982 [64]). *A directed graph is k -edge-connected if and only if it can be obtained from a single node by iteratively applying the following two operations:*

- (i) *add a new edge (possibly a loop),*
- (ii) *subdivide k existing edges and identify the subdividing nodes with a single node z .*

In this theorem and in Theorem 1.25 as well, operation (ii) is called **pinching** k edges with z . By pinching 0 edges we mean the addition of a node. Note that using Theorem 1.26, Theorem 1.25 can easily be derived from Theorem 1.27.

In the proof of Theorem 1.27, an intrinsic tool is another deep theorem of Mader on directed splitting off. Similarly to undirected graphs, in a digraph $G = (V, A)$, **splitting off** edges $e = xz$ and $f = zy$ means the operation of deleting e and f and adding the new edge xy . If $\rho(z) = \delta(z)$, a **complete splitting at z** is a sequence of splitting off operations of all edges incident to z and finally removing z . We say that a digraph $D = (U + z, A)$ is **k -edge-connected in U** if there are k -edge-disjoint directed paths between any two nodes in U .

Theorem 1.28 (Mader, 1982 [64]). *Let $D = (U + z, A)$ be a digraph which is k -edge-connected in U and $\rho(z) = \delta(z)$. Then there exists a complete splitting at z resulting in a k -edge-connected digraph.*

From Theorem 1.27 one may also derive the constructive characterization of rooted k -edge-connected digraphs (see e.g. [24]). A digraph $D = (V, A)$ is called **rooted k -edge-connected** if for a node $r_0 \in V$, there are k -edge-disjoint paths from r_0 to every node in $V - r_0$. Clearly, this is equivalent to $\rho(X) \geq k$ for every $X \subseteq V - r_0$.

Theorem 1.29. *A directed graph $D = (V, A)$ is rooted k -edge-connected with a root $r_0 \in V$ if and only if it can be obtained from the single node r_0 by iteratively applying the following two operations.*

(i) *add a new edge (possibly a loop),*

(ii) *pinch $0 \leq j \leq k - 1$ edges with a new node z and add $k - j$ new edges with head z .*

From this theorem, one may easily derive Edmonds' classical theorem on disjoint arborescences:

Theorem 1.30 (Edmonds, 1973 [17]). *A directed graph $D = (V, A)$ contains k edge disjoint spanning arborescences with root $r_0 \in V$ if and only if it is rooted k -edge-connected with root r_0 .*

Similarly to Theorem 1.26, rooted k -edge-connectivity of digraphs also has an undirected counterpart. An undirected graph is called **k -partition-connected** if for any partition of the node set into $t \geq 2$ classes, there are at least $k(t - 1)$ edges between different classes of the partition. Note that this is a property stronger than k -edge-connectivity.

Theorem 1.31 (Frank, 1980 [22]). *An undirected graph $G = (V, E)$ has a rooted k -edge-connected orientation with a root $r_0 \in V$ if and only if it is k -partition-connected.*

From this orientation theorem and Edmonds' theorem we can easily obtain Tutte's theorem:

Theorem 1.32 (Tutte, 1961 [71]). *An undirected graph contains k edge-disjoint spanning trees if and only if it is k -partition-connected.*

We can also derive the following characterization from Theorems 1.29 and 1.31:

Theorem 1.33. *An undirected graph is k -partition-connected if and only if it can be obtained from a single node by iteratively applying the following two operations.*

(i) *add a new edge,*

(ii) *pinch $0 \leq j \leq k - 1$ edges with a new node z and add $k - j$ new edges incident to z .*

(k, ℓ) -edge-connectivity is a natural common generalization of k -edge-connectivity and rooted k -edge-connectivity of digraphs. We say that $D = (V, A)$ is **(k, ℓ) -edge-connected** for some integers $0 \leq \ell \leq k$ and a root node $r_0 \in V$, if for each node $v \neq r_0$, there exist k edge-disjoint paths from r_0 to v and ℓ edge-disjoint paths from v to r_0 . Note that (k, k) -edge-connectivity coincides with k -edge-connectivity, while $(k, 0)$ -edge-connectivity means rooted k -edge-connectivity. Theorem 1.28 can also be extended to (k, ℓ) -edge-connectivity. We say that the digraph $D = (U + z, A)$ is **(k, ℓ) -edge-connected in U** for a root node $r_0 \in U$, if for every node $v \in U - r_0$ there are k -edge-disjoint paths from r_0 to v and ℓ edge-disjoint paths from v to r_0 in D .

Theorem 1.34 (Frank, 1999 [24]). *Let $D = (U + z, A)$ be a digraph (k, ℓ) -edge-connected in U and $\rho(z) = \delta(z)$. Then there exists a complete splitting at z resulting in a (k, ℓ) -edge-connected graph.*

Let us mention that this is still only a special case of Theorem 6.19, which can also be derived from Theorem 1.1. The analogous concept for undirected graphs is the following. An undirected graph is called (k, ℓ) -partition connected if for any partition of the nodes into $t \geq 2$ classes, there are at least $k(t - 1) + \ell$ edges connecting distinct classes. The link between these concepts is the following generalization of Theorem 1.31.

Theorem 1.35 (Frank, 1980 [22]). *For integers $0 \leq \ell \leq k$, an undirected graph G has a (k, ℓ) -edge-connected orientation if and only if G is (k, ℓ) -partition connected.*

Hence a natural problem arising is the constructive characterization of (k, ℓ) -edge-connected graphs, solved in Theorem 1.47 of this thesis. Based on Theorem 1.35, this will immediately give a constructive characterization of (k, ℓ) -partition-connected graphs. Besides $\ell = 0$ and $\ell = k$, the following special cases of Theorem 1.47 were known beforehand. $\ell = 1$ was shown by Frank and Szegő [34], and the case $\ell = k - 1$ was proved by Frank and Király [33]. Let us exhibit a nice application of the latter case.

An important open question is the following. Given an undirected graph $G = (V, E)$ and a subset of nodes $T \subseteq V$, we call an orientation of G **T -odd** if the nodes with odd in-degree are exactly those in T . The question is: *for a given node set T , decide whether there exists a strongly connected T -odd orientation*. A trivial necessary condition is that $|T| + |E|$ should be even, but no necessary and sufficient condition is known. However, we may ask whether there

is a strongly connected T -odd orientation for every $T \subseteq V$ with $|T| + |E|$ even. This question can be answered not only for strongly connectedness but for higher connectivity as well:

Theorem 1.36 (Frank and Király, 2002 [33]). *For an undirected graph $G = (V, E)$, the following three properties are equivalent:*

- (1) G has a k -edge-connected T -odd orientation for every $T \subseteq V$ with $|T| + |E|$ even.
- (2) G is $(k + 1, k)$ -partition connected.
- (3) G can be built up from a single node by a sequence of (i) adding new edges, and (ii) pinching k existing edges with a new node z and adding a new edge from an existing node to z .

At first sight it is neither clear if property (1) is in NP, nor if it is in co-NP. Property (2) gives a co-NP certificate: given a deficient partition, it is easy to construct a T not admitting a k -edge-connected T -odd orientation. On the other hand, (3) gives an NP-certificate: using the construction sequence, it is easy to find a good T -odd orientation for any T with $|T| + |E|$ odd. This application has motivated the investigation of $(k, k - 1)$ -partition-connected graphs.

1.5 Overview of the main results

Chapters 2 and 3 are devoted to the directed and undirected connectivity augmentation problems are closely related and we outline them side-by-side. Afterwards, the subsequent three chapters will be discussed separately.

1.5.1 Augmenting directed and undirected connectivity by one

For directed connectivity augmentation by one, the size of an optimal augmenting edge set is given in Theorem 1.5. Let us now give a min-max formula for undirected connectivity augmentation by one, which was conjectured by Frank and Jordán [30] in 1994. The basic object analogous to strict one-way pairs will be clumps, a notion corresponding to tight node cuts.

In the $(k-1)$ -connected graph $G = (V, E)$, a subpartition $X = (X_1, \dots, X_t)$ of V with $t \geq 2$ is called a **clump** if $|V - \bigcup X_i| = k-1$ and $d(X_i, X_j) = 0$ for any $i \neq j$. The sets X_i are called the **pieces** of X while $|X|$ denotes t , the number of pieces. If $t = 2$ then X is a **small clump**, while for $t \geq 3$ it is a **large clump**. (The set $V - \bigcup X_i$ is often called **separator** in the literature, and **shredder** if $t \geq 3$.) An edge $uv \in \binom{V}{2}$ **connects** X if u and v lie in different pieces of X . Two clumps are said to be **independent** if there is no edge $uv \in \binom{V}{2}$ connecting both.

A **bush** \mathcal{B} is a set of pairwise distinct small clumps, so that each edge in $\binom{V}{2}$ connects at most two of them. A **shrub** is a set consisting of pairwise independent (possibly large) clumps. For a bush \mathcal{B} let $def(\mathcal{B}) = \left\lceil \frac{|\mathcal{B}|}{2} \right\rceil$, and for a shrub \mathcal{S} let $def(\mathcal{S}) = \sum_{K \in \mathcal{S}} (|K| - 1)$.

A **grove** is a set consisting of some (possibly zero) bushes and one (possibly empty) shrub, so that the clumps belonging to different bushes are independent, and a clump belonging to a bush is independent from all clumps belonging to the shrub. For a grove Π consisting of the shrub \mathcal{B}_0 and bushes $\mathcal{B}_1, \dots, \mathcal{B}_\ell$, let $def(\Pi) = \sum_i def(\mathcal{B}_i)$. For a $(k-1)$ -connected graph $G = (V, E)$, let $\tau(G)$ denote the minimum number of edges whose addition makes G k -connected, and let $\nu(G)$ denote the maximum value of $def(\Pi)$ over all groves Π .

Theorem 1.37. *For a $(k-1)$ -connected graph $G = (V, E)$ with $|V| \geq k+1$, $\nu(G) = \tau(G)$.*

The theorem is illustrated in Figure 1.1. Both Chapters 2 and 3 contain algorithms using a dual oracle. Assume we are given a subroutine for determining the optimum value $\nu = \tau$ along an optimal dual structure. Based on this, the following simple algorithm gives a primal optimal solution. For an undirected graph $G = (V, E)$, let $J = \binom{V}{2} - E$ denote the edge set of the complement graph of G . Let us start with computing $\nu(G)$. In each step, choose an $e \in J$, and remove e from J . If $\nu(G+e) = \nu(G) - 1$, then add the edge e to G , otherwise keep the same G . The same algorithm works for a directed graph $D = (V, A)$, starting with $J = V^2 - A$. Note that Theorem 1.37 (in the directed case, Theorem 1.5) ensures the existence of an edge e with $\nu(G+e) = \nu(G) - 1$.




Figure 1.1: Let G be the graph in the figure with the addition of a complete bipartite graph between V_A and V_B and let $k = 8$. G is 7-connected, and it can be made 8-connected by the addition of the edge set $\{a_1a_3, a_2a_4, a_3a_5, b_3b_4, b_4b_5\}$. Two clumps $(\{a_1\}, \{a_3, a_4\})$ and $(\{b_3\}, \{b_4\}, \{b_5\})$ are shown on the figure. A grove Π with $\text{def}(\Pi) = 5$ consists of the shrub \mathcal{B}_0 and the bush \mathcal{B}_1 with $\mathcal{B}_0 = \{(\{b_3\}, \{b_4\}, \{b_5\})\}$, and $\mathcal{B}_1 = \{(\{a_1\}, \{a_3, a_4\}), (\{a_2\}, \{a_4, a_5\}), (\{a_3\}, \{a_5, a_1\}), (\{a_4\}, \{a_1, a_2\}), (\{a_5\}, \{a_2, a_3\})\}$.

For strict one-way pairs, we have already defined the notion of independence and crossing families; these can be naturally extended to clumps. A major difference is that no natural partial order may be defined on clumps, however, nestedness can be introduced as a notion analogous to comparability. In both cases, a cross-free system is a special class of crossing families of pairs (resp. clumps) so that any two members are either independent or comparable (resp. nested). A key notion is skeleton: a cross-free system maximal for containment.

Theorems 2.1 and 3.12 state that the maximum dual value over the members of a skeleton is the same as over all strict one-way pairs (resp. clumps). Once having a skeleton, we will be able to determine the dual optimum value relatively easily. In the directed case, Dilworth's theorem on the maximum size of an antichain in a poset gives the dual optimum. For the undirected case, instead of Dilworth's theorem we use Fleiner's theorem [20] on covering symmetric posets by symmetric chains. This may be seen as a common generalization of Dilworth's theorem and the Berge-Tutte theorem on the maximum size of a matching in a graph. While Dilworth's theorem can be derived from the König-Hall theorem on finding a maximum matching in bipartite graphs, Fleiner's theorem may be itself deduced from the Berge-Tutte theorem. The relation between directed and undirected connectivity augmentation is somewhat analogous, concerning both the complexity of the min-max formulae and the difficulty of the proofs.

Two proofs will be presented for Theorem 2.1. In Section 2.1 we give a simple, direct proof, while Section 2.2 contains a more complicated one. In the latter one, we start from an edge set F covering all strict one-way pairs in a given skeleton. By **flipping** two edges $xy, uv \in F$

mean replacing F by $F' = F - \{xy, uv\} + \{xv, uy\}$. (We use this definition both in directed and undirected graphs.) We prove that by a sequence of such operations we can arrive from F to a covering of all strict one-way pairs, that is, an augmenting edge set for D . The advantage of this latter proof is fourfold. First, it gives a proof not only for Theorem 2.1 but also for Theorem 1.5. Second, it enables us to construct an algorithm that calls the dual oracle only once. Third, it extends to node-induced cost functions as well. Finally, the greatest advantage is that the argument carries over with only minor changes to the undirected case.

In contrast to the astonishingly simple original proof of Theorem 1.1 and the direct proof of Theorem 2.1 in Section 2.1, the only method known so far for proving Theorems 1.37 and 3.12 is the adaptation of the argument of Section 2.2. However, I strongly believe that developing simpler proofs should be possible. In fact, Theorems 1.37 should be seen as a starting point rather than a final achievement in the area. I insist that it should be generalizable not only for general connectivity augmentation, but it should also admit a general abstract form analogous to Theorem 1.1. This generalization should include, among others, rooted connectivity augmentation and K_{tt} -free t -matchings (see [8]).

The main algorithmic task for the dual oracle is constructing a skeleton. Although any maximal cross-free system of strict one-way pairs (resp. clumps) suits, it is not trivial to find one since the number of strict one-way pairs and clumps may be exponentially large. To tackle this problem, the notion of stability of cross-free systems is defined in both cases. For stable cross-free systems, it will be fairly easy to determine whether they are skeletons, and if not, we will be able to extend them preserving stability. Although the structural properties are quite analogous, the argument in the undirected case will be significantly more complicated.

1.5.2 General connectivity augmentation

The approach in Chapter 4 for directed connectivity augmentation is completely different from the one in Chapter 2. This result is an extension of the previous work of Benczúr [4] on augmenting directed connectivity by one. The present result is applicable not only to directed connectivity augmentation, but gives a new, algorithmic proof of Theorem 1.1 (similarly, the result in [4] also worked for the more general Theorem 1.2).

Dilworth's theorem plays an important role in Chapter 2 since it is used for determining the maximum number of pairwise independent strict one-way pairs in a skeleton. Although not applied directly, it serves as a starting point and motivation for the current approach. We give a more general algorithm that resembles the version of Dilworth's algorithm described in [21]. The main theorem (Theorem 1.40) is an equivalent reformulation of Theorem 1.1 in terms of posets, for the problem of covering a certain type of weighted poset by a minimum number of intervals.

Definition 1.38. Consider a poset (\mathcal{P}, \preceq) . We say that for a minimal element m and a maximal

element M , the set $\{z : m \preceq z \preceq M\}$ is the **interval** $[m, M]$. Let $x, y \in \mathcal{P}$ be called **dependent** if there exists an interval $[m, M]$ with $x, y \in [m, M]$; otherwise they are called **independent**. We say that (\mathcal{P}, \preceq) satisfies the **strong interval property** if the following hold:

- (i) For all dependent $x, y \in \mathcal{P}$ the operations $x \vee y = \min\{z : z \succeq x, z \succeq y\}$ and $x \wedge y = \max\{z : z \preceq x, z \preceq y\}$ are uniquely defined.
- (ii) For every interval $[m, M]$,

$$x \wedge y \in [m, M] \text{ implies } x \in [m, M] \text{ or } y \in [m, M],$$

and the same holds with $x \wedge y$ replaced by $x \vee y$.

The notion of a **positively crossing supermodular function** p on such a poset is analogous to the one on set pairs: for all dependent x and y with $p(x) > 0$ and $p(y) > 0$ we require

$$p(x) + p(y) \leq p(x \wedge y) + p(x \vee y).$$

Consider a multiset of intervals \mathcal{I} . We say that \mathcal{I} **covers** the function p or \mathcal{I} is a **cover** of p if for every x , at least $p(x)$ intervals in \mathcal{I} contain x . An element v is called **tight** if contained in exactly $p(x)$ intervals in \mathcal{I} .

Given the notion of the cover problem for a poset with the strong interval property, we next show its equivalence to Theorem 1.1. We start with describing the correspondence between set pairs and poset elements as illustrated in Fig. 1.2. Property (ii) in the definition can be seen as the abstraction of the simple Lemma 2.2 for set pairs.




Figure 1.2: The correspondence between set pairs and poset elements. The four pairs on the left side can be covered by one edge, and the corresponding four elements are contained in one interval.

Claim 1.39. *The poset of set pairs (\mathcal{S}, \preceq) with the operations \wedge, \vee satisfy Definition 1.38. The set of intervals of this poset is $\{I_{uv} : uv \in V^2\}$, where $I_{uv} = \{K \in \mathcal{S} : uv \text{ covers } K\}$.*

Let us now formulate our theorem, which is an analogue of Theorem 1.1 for posets.

Theorem 1.40. *For a poset (\mathcal{P}, \preceq) with the strong interval property and a positively crossing supermodular function p , the minimum number of intervals covering p is equal to the maximum of the sum of p values of pairwise independent elements of \mathcal{P} .*

Using Claim 1.39, this theorem implies Theorem 1.1. We will show that the reverse is also true: this theorem can also be derived from Theorem 1.1.

Our algorithm uses a primal-dual scheme for finding covers of the poset. For an initial (possibly greedy) cover the algorithm searches for witnesses for the necessity of each element in the cover. If any two (weighted) witnesses are independent, the solution is optimal. As long as this is not the case, the witnesses are gradually exchanged by smaller ones. Each witness change defines an appropriate change in the solution; these changes are finally unwound in a shortest path manner to obtain a solution of size one less.

The algorithm itself is not very complicated (yet far from simple); however, the proof of correctness is technically quite involved. When applying it to concrete problems such as directed connectivity augmentation, we have to be careful since the size of the poset is typically exponential. The basic steps of the algorithm involve operations as finding the (unique) maximal tight element of an interval in a certain cover. In Section 4.2 we show that for directed connectivity augmentation, such oracle calls can be implemented via maximum flow computations.

The algorithm is pseudopolynomial as the size of the initial cover depends on the maximum value of p , and the size of the cover is increased by only one in each step. Of course, for connectivity augmentation this does not matter as the maximum value of p is at most $k \leq |V| - 1$; however, for ST -edge-connectivity augmentation, p may take arbitrarily large values.³ Hence developing a strongly polynomial or at least a polynomial time algorithm is still an important challenge.

1.5.3 Local edge-connectivity augmentation

Chapter 5 commences with new proofs of Theorems 1.17 and 1.22. Then we turn to the problem of partition-constrained local edge-connectivity augmentation (PCLECA). First, an approximation algorithm is presented for finding an augmenting edge set of \mathcal{Q} -legal edges of size at most the optimum plus r_{\max} , the largest connectivity requirement. Then, for the bipartite case (that is, if \mathcal{Q} consists of two classes) we formulate a conjecture on the minimum size of a \mathcal{Q} -legal augmenting edge set. We only give a partial proof of this conjecture, already extremely complicated. The completion of the proof and the extension to arbitrary number of partition classes is left for future research.

³Recall that the definition of k -node-connectivity also imposed $k \leq |V| - 1$; no similar restrictions exist for edge-connectivity and thus we may have an arbitrary requirement k independently from $|V|$.

To our best knowledge, all undirected edge-connectivity augmentation results discussed in Section 1.3 among their extensions (see e.g. the thesis of Bernáth [9]) were proved via splitting off techniques. We break this tradition by applying the alternative technique of edge-flippings. Consider a covering problem of a set function p and a degree-prescription m . Vaguely speaking, we want to find an m -prescribed edge-set F covering p “as much as possible”. For an m -prescribed edge set F , let us define the function

$$q_F(X) = p(X) - d_F(X).$$

Let $\nu_F = \max_{X \subseteq V} q_F(X)$. Note that F covers p if and only if $\nu_F = 0$. We will be interested in m -prescribed edge sets minimizing ν_F . Let

$$\mathcal{F}_F := \{X \subseteq V \mid q_F(X) = \nu_F \text{ and } \forall U \subsetneq X : q_F(U) < \nu_F\}$$

Let us define a partial order \preceq on the m -prescribed edge sets: $F' \prec F$ if $\nu_{F'} < \nu_F$, or $\nu_{F'} = \nu_F$ and $|\mathcal{F}_{F'}| < |\mathcal{F}_F|$. We are going to focus on \preceq -minimal m -prescribed edge sets. What we really use is the local optimality of such an F : with a small elementary change, we cannot get an F' from F with $F' \prec F$.

Recall that for two edges $xy, uv \in F$, by **flipping** (xy, uv) we mean replacing F by $F' = F - \{xy, uv\} + \{xv, uy\}$. In most proofs, it will be enough to assert that from a given F , we cannot get an $F' \prec F$ by a single flipping. Consequently, a local search algorithm can be applied for finding an optimal solution, given that we have oracles for determining the values ν_F and $|\mathcal{F}_F|$.

It turns out that for Theorems 1.17 and 1.22, a quite weak property of the demand function p almost suffices. p is called **symmetric positively skew supemodular** (abbreviated **SPSS**) if p is a nonnegative integer-valued function on the ground set V ; $p(X) = p(V - X)$ for every $V \subseteq X$, and for every pair $X, Y \subseteq V$ with $p(X), p(Y) > 0$, at least one of the following inequalities hold:

$$p(X) + p(Y) \leq p(X \cap Y) + p(X \cup Y), \tag{1.7a}$$

$$p(X) + p(Y) \leq p(X - Y) + p(Y - X) \tag{1.7b}$$

One basic example of such a function is $p(X) = (R(X) - d(X))^+$ for $R(X)$ defined by a local edge-connectivity requirement, while the other example is a symmetric and positively crossing supermodular function. Although covering an arbitrary SPSS-function is NP-complete (see [54]), it is easy to find an edge set almost covering p . Namely, we prove the following.

Theorem 1.41. *Let p be an SPSS-function and m a degree-prescription so that (1.4) holds. For a \preceq -minimal m -prescribed edge set F , $\nu_F \leq 1$ holds, or equivalently, $d_F(X) \geq p(X) - 1$ for every $X \subseteq V$.*

Therefore, in both Theorems 1.17 and 1.22 it will be enough to focus on the case $\nu = 1$. For this, stronger properties of the particular function p are needed. Theorem 1.41 is a folklore result, appearing in the thesis of Cosh [15], in the papers of Nutov [67] and Bernáth and Király [12].

Edge-flipping is a classical technique for degree-prescribed problems: see for example, Hakimi’s paper [39] from 1962 or Edmonds’ result [16] from 1964. For digraphs, Frank and Z. Király [33] applied a similar technique to give a new proof of Theorem 6.19, a generalization of Theorem 1.28 on directed splitting off.

For Theorem 1.17, we do not claim that edge-flipping leads to a much easier proof. For Theorem 1.22, the two proofs known by the author are the original one by Benczúr and Frank [6], and a recent, significantly simpler one by Bernáth [10]. Let us take a degree-prescription m satisfying (1.4) and add a new node z connected to each node v by $m(v)$ parallel edges. In the case of Theorem 1.17, an arbitrary sequence of legal splittings was feasible, however, this does not apply for Theorem 1.22. Benczúr and Frank show the existence of “good” pair of splittable edges, nevertheless, tremendous technical effort is required to find such a pair. If we cannot remove all edges incident to z this way, then a p -full partition can be exhibited, showing that (1.6) did not hold originally. On the contrary, Bernáth proceeds by splitting arbitrary feasible pairs of edges as long as possible. The drawback of this method is that we are not finished in the case when no complete splitting exists. It needs to be checked whether we can obtain a better situation by undoing a previous splitting off, similarly to the method of Bang-Jensen et al. [2] as sketched after Theorem 1.23.

In contrast, our proof of Theorem 1.22 is quite analogous to that of Theorem 1.17. Consider a degree-prescription m satisfying (1.4) and choose an m -prescribed edge set F so that we cannot get an F' with $F' \prec F$ by performing a single edge flipping. In both cases, such an F is optimal: in Theorem 1.17 we can deduce $\nu = 0$ while in Theorem 1.22 $\nu \leq 1$, and if $\nu = 1$ then (1.6) does not hold. The proof of $\nu \leq 1$ is provided by the same Theorem 1.41 in both cases.

My main motivation for applying edge-flippings in the context of undirected covering problems was the hope that it could be more suitable for the PCLECA problem. Splitting off with the aforementioned technique of undoing splittings is also a natural way to attack this problem, and I also started this way. The main difficulty is that, in contrast to global edge-connectivity, undoing a single splitting off is insufficient. I conjecture that undoing two should be enough; however, at a certain point the analysis becomes severely complicated. I think that edge-flipping is more appropriate to tackle this problem. Unfortunately, I could neither complete the proof with this method, however, I think that the partial results might be of some value.

For both augmentation Theorems 1.15 and 1.21, we had the degree-prescribed versions Theorems 1.17 and 1.22. Let us now formulate the degree-prescribed version of the PCLECA problem. For some integer $t \geq 2$, let us be given degree sequences $m_1, \dots, m_t : V \rightarrow \mathbb{Z}_+$, and let

$m = \sum_{i=1}^t m_i$. $\vec{m} = (m_1, \dots, m_t)$ is called a **legal degree-prescription** if $m(V)$ is even and

$$m_i(V) \leq \frac{m(V)}{2} \text{ for } i = 1, \dots, t. \quad (1.8)$$

The integers $\{1, \dots, t\}$ will be called **colours**. Notice that for $t = 2$, (1.8) is equivalent to $m_1(V) = m_2(V) = \frac{1}{2}m(V)$. Consider a pair (F, φ) consisting of an edge set F equipped with a mapping φ . This maps the endnodes of the edges in F to the set of colours so that for $xy \in F$, $\varphi(xy, x) \neq \varphi(xy, y)$. An edge $xy \in F$ with $\varphi(xy, x) = i$ and $\varphi(xy, y) = j$ is called an ij -edge.⁴ (F, φ) is called an \vec{m} -**prescribed legal edge set**⁵ if

$$|\{xy : \varphi(xy, x) = i\}| = m_i(x) \text{ for } x \in V, i = 1, \dots, t. \quad (1.9)$$

It can be seen easily that if \vec{m} is a legal degree-prescription, then there exists an \vec{m} -prescribed legal edge set. Edge-flippings can be naturally defined with taking the mapping φ also into account. The difference is that for $xy, uv \in F$, flipping (xy, uv) is possible only if $\varphi(xy, x) \neq \varphi(uv, v)$, $\varphi(xy, y) \neq \varphi(uv, u)$. Nevertheless, at least one of (xy, uv) and (xy, vu) can be flipped.

Given a partition $\mathcal{Q} = \{Q_1, \dots, Q_t\}$ and a degree-prescription $m : V \rightarrow \mathbb{Z}_+$, we may define $m_i(v) = m(v)$ if $v \in Q_i$ and $m_i(v) = 0$ otherwise. (Note that this is not always a legal degree-prescription as (1.8) is not necessarily satisfied.) The model above is slightly more general since we allow $m_i(v) = m_j(v) > 0$ for $i \neq j$. We advise the reader to keep this example in mind in the sequel; note that here φ is uniquely defined by the partition \mathcal{Q} .

Given the connectivity requirement function r , we are interested in coverings of the function $p(X) = (R(X) - d(X))^+$ by \vec{m} -prescribed legal edge sets. (1.4) is a necessary, but not sufficient condition. For a legal degree-prescription \vec{m} satisfying (1.4), we will be interested in minimizing ν_F over \vec{m} -prescribed legal edge sets. The first, relatively simple result we prove in Section 5.2.1 is the following.

Theorem 1.42. *Given r and a legal degree-prescription \vec{m} satisfying (1.4), consider an \vec{m} -prescribed \preceq -minimal F . If $\nu_F > 0$ then $|\mathcal{F}_F| = 2$.*

This theorem will enable us to construct a simple approximation algorithm for the PCLECA problem in Section 5.2.2 with an additive term r_{\max} .

Theorem 1.43. *Assume we are given a graph $G = (V, E)$, a partition \mathcal{Q} of the nodes and a connectivity requirement r so that G contains no marginal sets. Then the minimum number of \mathcal{Q} -legal edges whose addition makes G r -edge-connected is at most $\Psi_{\mathcal{Q}}(G) + r_{\max}$.*

⁴Denoting the same edge by xy or yx has different meanings, as the one is an ij -edge while the other a ji -edge. For $t = 2$, we could also represent F by directed edges.

⁵We will often omit φ and refer only to F as an \vec{m} -prescribed legal edge set. Nevertheless, φ is always tacitly included. For example, we speak of ij -edges in F .

Recently, a weaker version of this theorem was also proved by Lau and Yung [58] (for two partition classes and $2r_{\max}$.)

For $t = 2$, we formulate conjectures on the optimum value of ν_F in the degree-prescribed problem and on the minimum size of a \mathcal{Q} -legal augmenting edge set in the augmentation problem. The dual structure is given by the next sophisticated definition.

Definition 1.44. Consider a partition $\mathcal{H} = \{X^*, Y^*, C_1, C_2, \dots, C_\ell\}$ of V . We say that \mathcal{H} forms a **hydra** with **heads** X^*, Y^* and **tentacles** C_i if

- (i) $d_G(C_i, C_j) = 0$ for every $1 \leq i < j \leq \ell$; and
- (ii) For any two disjoint index sets $\emptyset \neq I, J \subseteq \{1, \dots, \ell\}$, (1.5a) holds with equality for $X^* \cup (\bigcup_{i \in I} C_i)$ and $X^* \cup (\bigcup_{j \in J} C_j)$, and also for $Y^* \cup (\bigcup_{i \in I} C_i)$ and $Y^* \cup (\bigcup_{j \in J} C_j)$.

Similarly to p -full partitions, although the definition contains exponentially many conditions, Theorem 5.23 will give an equivalent characterization in terms of the values of r between different classes of \mathcal{H} . This also yields an efficient method to decide whether a partition forms a hydra.

Given a requirement function r , a legal degree-prescription $\vec{m} = (m_1, m_2, \dots, m_t)$ and $1 \leq h \leq t$, we call a tentacle C_i **h -odd** if $p(C_i \cup X^*) - p(X^*) + m_h(C_i)$ is odd.⁶ Let χ_h denote the number of h -odd tentacles. Let us define

$$\tau_h(G, r, \vec{m}, \mathcal{H}) = \frac{1}{2} \left(\chi_h + p(X^*) + p(Y^*) - m(V) + m_h(\bigcup C_i) \right).$$

Let

$$\tau(G, r, \vec{m}) = \max\{0, \max_{h=1}^t \tau_h(G, r, \vec{m}, \mathcal{H}) : \mathcal{H} \text{ is a hydra}\}$$

The conjecture on the degree-prescribed version of the PCLECA problem is as follows.

Conjecture 1.45. *Let us be given a graph $G = (V, E)$ with a connectivity requirement function r so that G contains no marginal sets. If $\vec{m} = (m_1, m_2)$ is a legal degree-prescription satisfying (1.4) and (F, φ) is a \preceq -minimal \vec{m} -prescribed legal edge sets, then $\nu_F = \tau(G, r, \vec{m})$.*

The corresponding conjecture for the augmentation problem is as follows. Let $\mathcal{Q} = \{Q_1, Q_2\}$ be the partition constraint. Let $\mathcal{H} = (X^*, Y^*, C_1, C_2, \dots, C_\ell)$ be a hydra, and for $h \in \{1, 2\}$, let \mathcal{Z} be an h -subpartition which is a refinement of $\{C_1, \dots, C_\ell\}$. (Recall that by an h -subpartition we mean a subpartition of Q_h .) The tentacle C_i is called **h -toxic** if

$$p(C_i \cup X^*) - p(X^*) + \sum (p(Z) : Z \in \mathcal{Z}, Z \subseteq C_i)$$

is odd. Let χ'_h denote the number of h -toxic tentacles. Let us define

$$\tau'_h(G, r, \mathcal{Z}, \mathcal{H}) = \frac{1}{2} (\chi'_h + p(X^*) + p(Y^*) + p(\mathcal{Z})).$$

⁶In Lemma 5.26 we shall prove that $p(C_i \cup X^*) - p(X^*) = -(p(C_i \cup Y^*) - p(Y^*))$, thus the role of X^* and Y^* is interchangeable.




Figure 1.3: Let $r(x, y) = 8$ and $r(u, v) = 3$ for any other pair. Let $Q_1 = \{x, y\}$ and $Q_2 = \{c_1, \dots, c_6\}$ be the partition classes. We have a hydra \mathcal{H} with $X^* = \{x\}$, $Y^* = \{y\}$, $C_i = \{c_i\}$ for $i = 1, \dots, 6$. Consider the degree-prescription $m_1(x) = m_1(y) = 3$, $m_2(c_i) = 1$ for $i = 1, \dots, 6$ and $m_j(u) = 0$ otherwise. All components C_i are 2-odd, $p(X^*) = p(Y^*) = 2$ and thus $\tau_2(G, r, \vec{m}, \mathcal{H}) = 2$. For the augmentation version, take the 2-subpartition \mathcal{Z} consisting of the singletons $\{c_i\}$. Then all C_i -s are 2-toxic, and $\tau'_2(G, r, \mathcal{Z}, \mathcal{H}) = 8$.

Let $\tau'(G, r, \mathcal{Q})$ denote the maximum of $\tau'_h(G, r, \mathcal{Z}, \mathcal{H})$ over all choices of h , \mathcal{H} and \mathcal{Z} as above. Recall that $\Psi_{\mathcal{Q}}(G)$ was defined in Section 1.3 as the maximum of $\alpha(G)$ and $\beta_j(G)$ for $j = 1, \dots, t$.

Conjecture 1.46. *Let us be given a graph $G = (V, E)$ with a connectivity requirement function r so that G contains no marginal sets and furthermore a partition $\mathcal{Q} = \{Q_1, Q_2\}$ of V . Then the minimum size of a \mathcal{Q} -legal augmenting edge set equals the maximum of $\Psi_{\mathcal{Q}}(G)$ and $\tau'(G, r, \mathcal{Q})$.*

C_4 -configurations are special hydra-bounds: consider a partition (A_1, A_2, C_1, C_2) of V and a h -partition \mathcal{Z} forming a C_4 -configuration, Then $\mathcal{H} = (X^*, Y^*, C_1, C_2)$ forms a hydra for $X^* = A_1$, $Y^* = A_2$ with both C_1 and C_2 being h -toxic; from the properties in the definition it follows that $\tau'_h(G, r, \mathcal{H}, \mathcal{Z}) = \Psi_{\mathcal{Q}}(G) + 1$.

It is already nontrivial that $\tau(G, r, \vec{m})$ and $\tau'(G, r, \mathcal{Q})$ are lower bounds on the optimum values: this will be proved in Section 5.2.4. In Section 5.3, we prove Theorem 5.30, a special case of Conjecture 1.45 under the assumptions that for the optimal F , $\nu_F \geq 2$ and $\bigcup \mathcal{F}_F = V$. The proof is quite technical. First, we extract structural properties from the assumption that we cannot get a better F' from F by performing a flipping or a “hexa-flipping”, a sequence of two edge flippings. This results in a set system containing a set “blocking” the edges of F in a certain sense. Afterwards, a complicated uncrossing method is applied to transform this set system into a laminar one, yielding an optimal hydra.

We think that this method should be extendable for proving Conjecture 1.45, however, the extreme level of complexity and the time and space limitations have forbidden us to give a complete proof. Finally, in Section 5.3.2, we sketch how Conjecture 1.46 could be derived from Conjecture 1.45. Also, we think that the conjectures could easily be extended to arbitrary number of partition classes, by adding another type of lower bound generalizing C_6 -configurations. In the global connectivity version [2], the main difficulties are already contained in the case

$t = 2$; we believe that the situation here should be similar.

1.5.4 Characterization of (k, ℓ) -edge-connected digraphs

The main result of this chapter is the following constructive characterization of (k, ℓ) -edge-connected digraphs, conjectured by András Frank ([34], Conjecture 5.6. and [28], Conjecture 5.1):

Theorem 1.47. *For $0 \leq \ell \leq k - 1$, a directed graph is (k, ℓ) -edge-connected with root $r_0 \in V$ if and only if it can be built up from the single node r_0 by the following two operations.*

- (i) *add a new edge,*
- (ii) *for some i with $\ell \leq i \leq k - 1$, pinch i existing edges with a new node z , and add $k - i$ new edges entering z and leaving existing nodes.*

We get the following corollary using Theorem 1.35:

Theorem 1.48. *For $0 \leq \ell \leq k - 1$, an undirected graph is (k, ℓ) -partition-connected if and only if it can be built up from a single node by the following two operations.*

- (i) *add a new edge,*
- (ii) *for some i with $\ell \leq i \leq k - 1$, pinch i existing edges with a new node z , and add $k - i$ new edges between z and some existing nodes.*

In Theorem 1.47, it is straightforward that all graphs constructed by operations (i) and (ii) are (k, ℓ) -edge-connected, the nontrivial part is the opposite direction. Removing an edge is the reverse of operation (i), hence we may focus our attention to minimally (k, ℓ) -edge-connected digraphs in the sense that removing any edge would destroy (k, ℓ) -edge-connectivity.

Let us sketch a proof of Theorem 1.27, which is a starting point of our argument (and corresponds to the special case $k = \ell$). If a digraph is not minimally k -edge-connected, we can leave an edge as the reverse of operation of step (i) and continue by induction. For minimally k -edge-connected digraphs, the existence of a node z having both in- and out-degree k can be proved. Then Mader's directed splitting theorem (Theorem 1.28) can be used since the reverse of operation (ii) is exactly a complete splitting at a node z with $\rho(z) = \delta(z)$. The case $\ell = 0$ (Theorem 1.29) can also be proved using an easy consequence of Theorem 1.28.

However, for the cases $\ell = 1$ and $\ell = k - 1$ of Theorem 1.47 we already need the stronger splitting result Theorem 1.34. The argument is also significantly more complicated for the following reason. For $\ell = k$ and $\ell = 0$, it was enough to find a node satisfying certain conditions on the in- and outdegrees, and one could always perform a complete splitting at such a node.

However, for $\ell = 1$ and $\ell = k - 1$ the conditions on the degrees do not suffice and a more thorough analysis of the structure of minimally (k, ℓ) -edge-connected graphs is needed.

Let us now sketch the proof for $\ell = k - 1$ by Frank and Király [33]. Consider a minimally $(k, k - 1)$ -edge-connected graph. A necessary condition for the reverse of operation (ii) to be applicable at node z is $\rho(z) = k$ and $\delta(z) = k - 1$. We call such nodes **special**. If for a special node z we manage to find an edge uz so that $D - uz$ is $(k, k - 1)$ -edge-connected in $U = V - z$, then Theorem 1.34 may be used for $D' = (U + z, A - uz)$, giving a $(k, k - 1)$ -edge-connected graph D'' on U . Then we can get D from D'' by applying step (ii) with pinching those $k - 1$ edges with z which were resulted by the splitting off and finally adding the edge uz .

However, not every special node z admits an edge uz as above (and it is already nontrivial to prove that a special node exists). We use an indirect argument: assume that every edge $xy \in A$ satisfies one of the following conditions. If y is special, then we assume that $D - xy$ is not $(k, k - 1)$ -edge-connected in $V - y$. If y is not special, we use that D is minimally $(k, k - 1)$ -edge-connected, and thus $D - xy$ is not $(k, k - 1)$ -edge-connected. One can define a notion of tight sets so that each edge will be “blocked” by a tight set. Then the uncrossing method may be used for these tight sets to derive a final contradiction.

The proof of Theorem 1.47 is motivated by this argument, but for general ℓ , severe difficulties arise. Starting from a minimally (k, ℓ) -edge-connected digraph, we call a node z **special** if $\ell \leq \delta(z) \leq k - 1$ and $\rho(z) = k$. This means that according to its in- and out-degree, it might be the result of operation (ii) in Theorem 1.47. We say that a subset F of edges entering a special node z is **locally admissible** at z if $G - F$ is (k, ℓ) -edge-connected in $V - z$ and $|F| \leq k - \delta(z)$. F is called **sufficient** at z if $|F| = k - \delta(z)$. Once a sufficient locally admissible F is found, Theorem 1.34 may be applied to $G - F$ and z and the proof finishes as for $\ell = k - 1$.

Thus our aim is to find a special node z and a sufficient locally admissible set F at z . It is easy to characterize the maximal size of a locally admissible set for a given special z , however, this size may be strictly smaller than $k - \delta(z)$. The main difficulty is handling the locally admissible sets belonging to different special nodes together. The notion of **globally admissible** edge sets in Definition 6.3 is introduced for this purpose. For a globally admissible edge set F and an arbitrary special node z , the subset $F_z \subseteq F$ of edges entering z is locally admissible at z . However, the converse is not true in the sense that the union of locally admissible edge sets belonging to different special nodes will not necessarily be globally admissible. We say that a globally admissible edge set F is **sufficient** if for some special z , F_z is sufficient; otherwise it is called **insufficient**. What we prove is the existence of a sufficient globally admissible edge set. Unfortunately, it is not true that every maximal globally admissible set is sufficient, as it will be shown by an example in Section 6.5.

Among other methods, splitting off techniques will be used also in the proof of the existence of a sufficient globally admissible set. However, even Theorem 1.34 turns out to be too weak for our

goals. Actually, Theorem 1.34 is a special case of Theorem 6.19 on covering positively crossing supermodular functions by a digraph. Theorem 6.20 is a further generalization presented in Section 6.3. It enables us to use a splitting operation preserving a property stronger than (k, ℓ) -edge-connectivity. The proof relies on edge flippings, used in an analogous manner as in Chapter 5 for undirected graphs.

The way we handle tight sets also differs from the standard uncrossing methods. A set is called **tight** with respect to a globally admissible set F if the inequality concerning this set in the definition of global admissibility holds with equality. As in the proof for $\ell = k - 1$, for a maximal F there is a tight set “blocking” each edge in $E - F$. However, it is not possible to apply the uncrossing method to arbitrary tight sets for an arbitrary globally admissible F . The intersection and union of two tight sets will be tight only under the assumption that F is maximal and insufficient. It turns out interestingly that under this assumption, some basic types of tight sets do not occur at all. This will be discussed in Section 6.4.

Contributions

Chapter 2 is based on a joint paper with András Frank [36], and Chapter 3 is based on the technical report [73]. The result of Chapter 4 is a joint work with András Benczúr in [74], while that of Chapter 6 is co-authored by Erika Renáta Kovács [56]. Chapter 5 contains unpublished material by the author.

Connections between the chapters

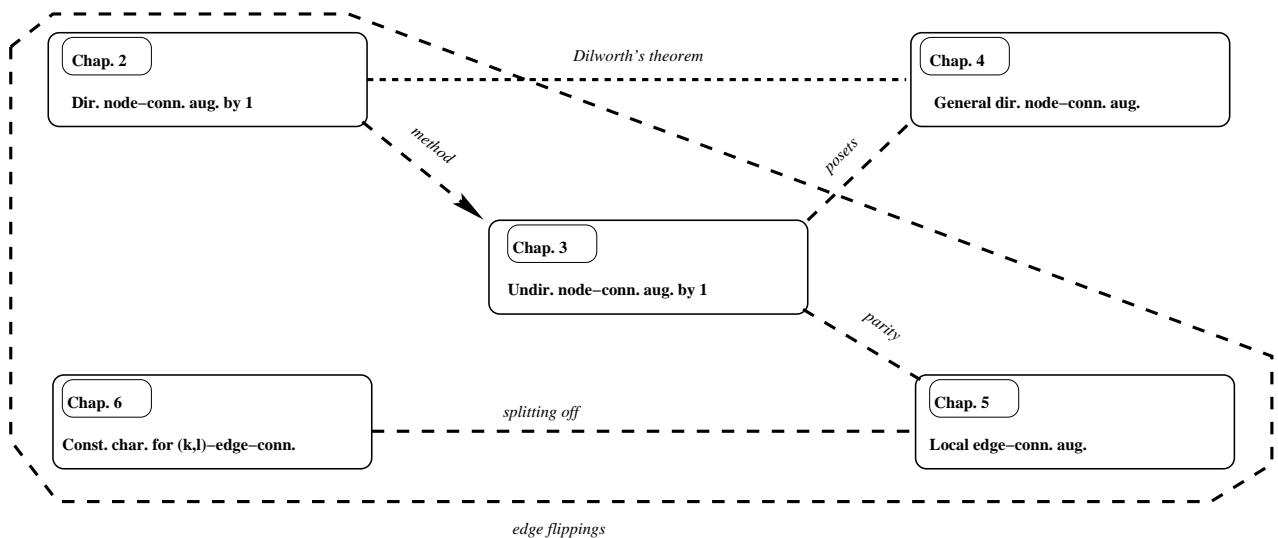


Figure 1.4: The hypergraph of interconnections.

At this point, the reader might have arrived to the conclusion that the thesis is rather a

compilation of scarcely related results with the author's person being the only common denominator. While we cannot completely refute such an opinion by exhibiting one common motif of the entire thesis, we tried to summarize some less transparent interconnections in Figure 1.4.

The most intimate relationship is indubitably the one between Chapters 2 and 3 on augmenting node-connectivity by one. We could adapt the main thoughts and structural elements of the proof of the directed case to the undirected case, albeit the min-max formulae being considerably different. In contrast, although Chapters 2 and 4 tackle the same problem, the methods do not have much in common. Nevertheless, we should mention Dilworth's theorem, which is applied in Chapter 2 directly and serves as a motivation for Chapter 4. As a connection between Chapters 3 and 4, we may exhibit the underlying poset structures. It is of key importance in both cases that we investigate the abstract poset properties of clumps and set pairs, respectively.

The occurrence of splitting off techniques in both Chapters 5 and 6 is quite natural: it is a fundamental and efficient method in edge-connectivity problems. Another method, edge-flipping is applied in various contexts in all but Chapter 4. On the one hand, it can be used as an alternative of splitting off: for example, in Chapter 5 we present new proofs of Theorems 1.17 and 1.22 using edge-flipping and we apply this technique for the PCLECA problem as well. The general directed covering result Theorem 6.20 is also proved via edge-flipping. On the other hand, in the completely different context of directed and undirected connectivity augmentation, the transformation of a cover of skeleton to a cover of all strict one-way pairs (resp. clumps) also relies on edge-flippings.

Chapters 3 and 5 share a somewhat odd common feature: parity is involved in both. It was known beforehand, that undirected node-connectivity augmentation has to do with parity, since it generalizes maximum matching. However, the emergence of parity might be surprising in the context of edge-connectivity. In Conjectures 1.45 and 1.46 there are certain odd components, resembling those in the Berge-Tutte formula. To the extent of my knowledge, parity has not been involved in such a way in previous edge-connectivity results. More interestingly, we conjecture that the optimum value described by these formulae can be found by a local search algorithm.

Chapter 2

Augmenting directed node-connectivity by one

In this chapter, we give an alternative proof and a combinatorial algorithm for Theorem 1.5, based on [36], a joint paper with András Frank. We will assume throughout the chapter that the digraph $D = (V, A)$ is $(k - 1)$ -connected. Let $\mathcal{O}^1 = \mathcal{O}_D^1$ denote the set of strict one-way pairs. Since we are now interested in strict one-way pairs only, we omit “strict” and use only “one-way pair” for the members of \mathcal{O}_1 . Some definitions and lemmas are formulated for set pairs; these are valid in the most general setting.

Let us start with some new notion. We have already introduced crossing families of set pairs in Section 1.1. A family $\mathcal{F} \subseteq \mathcal{S}$ is called **cross-free** if any two members of \mathcal{F} are either independent or comparable. Note that, somewhat confusingly, every cross-free family is crossing. For a set pair $K \in \mathcal{F}$, let $\mathcal{F} \div K$ denote the members of \mathcal{F} not crossing K . Similarly, for a subset $\mathcal{K} \subseteq \mathcal{F}$ let $\mathcal{F} \div \mathcal{K}$ denote the set of set pairs in \mathcal{F} crossing no element of \mathcal{K} . Let us call a cross-free subset $\mathcal{F} \subseteq \mathcal{O}^1$ a **skeleton** if $\mathcal{O}^1 \div \mathcal{F} = \mathcal{F}$. Equivalently, \mathcal{F} is a maximal cross-free subset of \mathcal{O}^1 .

In Section 2.1, we give the description of the Dual Oracle, a subroutine for determining $\nu(\mathcal{O}^1)$. In Section 2.1.2 we analyze the oracle and the first algorithm, which relies on this oracle. In Section 2.2, we give a new proof for Theorem 1.6, and sketch a second algorithm. For this algorithm, we present only the main ideas, and omit the technical details which can be done similarly as for the first algorithm.

2.1 The Dual Oracle

The following theorem is the essence of the Dual Oracle.

Theorem 2.1. *For a skeleton $\mathcal{K} \subseteq \mathcal{O}^1$ the maximum number of pairwise independent one-way pairs is equal in \mathcal{K} and \mathcal{O}^1 , that is, $\nu(\mathcal{K}) = \nu(\mathcal{O}^1) = \nu(D)$.*

Clearly, $\nu(\mathcal{K}) \leq \nu(\mathcal{O}^1)$ for every $\mathcal{K} \subseteq \mathcal{O}^1$. The advantage of a cross-free system is that we can easily determine the maximum number of pairwise independent one-way pairs. This is due to the fact that whenever it contains two dependent one-way pairs, they are comparable. Thus considering the partially ordered set (\mathcal{K}, \preceq) an antichain consists of pairwise independent pairs. A maximum antichain in a poset can be easily found by an algorithm for Dilworth's theorem stating the equality of the size of a minimum chain cover and a maximum antichain (see e.g. [69, Vol A., pp. 217-236]). In order to prove Theorem 2.1, we need some elementary propositions.

Lemma 2.2. *Let $M, N \in \mathcal{S}$ be two dependent set pairs. If an edge $xy \in V^2$ covers $M \wedge N$ or $M \vee N$, then it covers at least one of M and N . If it covers both $M \wedge N$ and $M \vee N$, then it covers both M and N . \square*

Claim 2.3. *Let $M, N \in \mathcal{O}^1$. $M^- \subseteq N^-$ implies $M \preceq N$, and $M^+ \subseteq N^+$ implies $M \succeq N$.*

Proof. For the first part, assume that $M \not\preceq N$, meaning that $M^+ \not\supseteq N^+$. Although M and N are not necessarily dependent ($M^+ \cap N^+ \neq \emptyset$ is not assumed), we may consider the set pair $L = (M^-, M^+ \cup N^+)$. This is a one-way pair, and since D is $(k-1)$ -connected, $s(L) \geq k-1$. However, M is a strict one-way pair, and since $M^+ \cup N^+ \supsetneq M^+$, we get $s(L) < s(M) = k-1$, a contradiction. The second part follows similarly. \square

Lemma 2.4. *For a crossing family \mathcal{F} and for any $K \in \mathcal{O}^1$, the subfamily $\mathcal{F} \div K$ is crossing.*

Proof. Let $\mathcal{F}' = \mathcal{F} \div K$ and let M and N be two crossing members of \mathcal{F}' . We have to prove that neither $M \vee N$ nor $M \wedge N$ crosses K .

First assume that K is comparable with both M and N . It is not possible that $M \preceq K \preceq N$ or $N \preceq K \preceq M$ as M and N are not comparable. Therefore either $K \preceq M, N$ or $K \succeq M, N$. In the first case, K is smaller than both $M \wedge N$ and $M \vee N$, while in the second case it is larger than both.

Second, assume that K is independent from both M and N . We claim that both $M \wedge N$ and $M \vee N$ are independent from K . Indeed, if an edge $xy \in V^2$ covered both K and $M \wedge N$ or $M \vee N$, then by Lemma 2.2, it would also cover M or N , a contradiction.

In the third case K is independent from one of M and N , say from M , and comparable with the other, N . If $K \preceq N$, then K and M can only be tail-disjoint, since $M^+ \cap N^+ \neq \emptyset$ and $K^+ \supseteq N^+$. Now $M \wedge N$ is also tail-disjoint from K , and $K \preceq M \vee N$. Similarly, if $K \succeq N$, then K and M should be head-disjoint, thus $M \vee N$ is head-disjoint from K , while $K \succeq M \wedge N$. \square

Lemma 2.5. *(i) Let L_1, L_2, L_3 be one-way pairs with L_1 and L_2 dependent, $L_1 \wedge L_2$ and L_3 also dependent, but L_2 and L_3 independent. Then $L_3^+ \cap (L_1^+ - L_2^+) \neq \emptyset$ and $L_1^- - L_2^- \subseteq L_3^-$. (ii) Let L_1, L_2, L_3 be one-way pairs with L_1 and L_2 dependent, $L_1 \vee L_2$ and L_3 also dependent, but L_2 and L_3 independent. Then $L_3^- \cap (L_1^- - L_2^-) \neq \emptyset$ and $L_1^+ - L_2^+ \subseteq L_3^+$.*

Proof. (i) The dependence of $L_1 \wedge L_2$ and L_3 implies $L_2^- \cap L_3^- \neq \emptyset$, so L_2 and L_3 can only be independent if $L_2^+ \cap L_3^+ = \emptyset$. The first part follows since $L_3^+ \cap (L_1^+ \cup L_2^+) \neq \emptyset$ because of the dependence of $L_1 \wedge L_2$ and L_3 . For the second part, consider the pair $N = (L_1 \wedge L_2) \vee L_3$. $N^+ = (L_1^+ \cup L_2^+) \cap L_3^+ = L_1^+ \cap L_3^+$, hence $N^+ \subseteq L_1^+$. By Claim 2.3, $N^- \supseteq L_1^-$, implying the claim. (ii) follows the same way, by exchanging the role of the tails and heads. \square

Now we are ready to prove Theorem 2.1. The proof is based on the following lemma:

Lemma 2.6. *For a crossing system \mathcal{F} and $K \in \mathcal{F}$ we have $\nu(\mathcal{F}) = \nu(\mathcal{F} \div K)$.*

First we show how Theorem 2.1 follows from Lemma 2.6. Let $\mathcal{K} = \{K_1, \dots, K_\ell\}$. First apply Lemma 2.6 for \mathcal{O}^1 and K_1 , then in the i th step for $\mathcal{O}^1 \div \{K_1, \dots, K_{i-1}\}$ and K_i . Note that $\mathcal{O}^1 \div \{K_1, \dots, K_{i-1}\}$ is a crossing system by applying inductively Lemma 2.4. Thus we have $\nu(\mathcal{O}^1) = \nu(\mathcal{O}^1 \div K_1) = \dots = \nu(\mathcal{O}^1 \div \mathcal{K})$, hence Theorem 2.1 follows by $\mathcal{O}^1 \div \mathcal{K} = \mathcal{K}$.

Proof of Lemma 2.6. Trivially, $\nu(\mathcal{F} \div K) \leq \nu(\mathcal{F})$. Consider a maximum independent subset \mathcal{L} of \mathcal{F} which has the most common members with $\mathcal{F} \div K$. For a contradiction, suppose that $\mathcal{L} \cap (\mathcal{F} \div K) < \nu(\mathcal{F})$, and choose an element $T \in \mathcal{L} - (\mathcal{F} \div K)$. By definition, T crosses K . We claim that either $(\mathcal{L} \setminus \{T\}) \cup \{T \wedge K\}$ or $(\mathcal{L} \setminus \{T\}) \cup \{T \vee K\}$ is independent. This leads to contradiction, since the new system intersects $\mathcal{F} \div K$ in a strictly larger subset than \mathcal{L} does.

Suppose that neither $(\mathcal{L} \setminus \{T\}) \cup \{T \wedge K\}$ nor $(\mathcal{L} \setminus \{T\}) \cup \{T \vee K\}$ is independent. Then there is an element $M \in \mathcal{L} \setminus \{T\}$ dependent from $T \wedge K$, and an other element $M' \in \mathcal{L} \setminus \{T\}$ dependent from $T \vee K$. If $M = M'$, then M is clearly dependent from L , a contradiction.

Assume now $M \neq M'$. The conditions of Lemma 2.5(i) hold for $L_1 = K$, $L_2 = T$ and $L_3 = M$, and the conditions of (ii) hold for $L_1 = K$, $L_2 = T$ and $L_3 = M'$. We claim that M and M' are dependent. Indeed, $K^- - L^-$ contains an element of $M^- \cap M'^-$, while $K^+ - L^+$ contains an element of $M^+ \cap M'^+$. \square

2.1.1 Constructing a skeleton

A straightforward approach to construct a skeleton of \mathcal{O}^1 would be a greedy method, that is, choose one-way pairs arbitrarily, as long as they do not cross any of the previously selected ones. The difficulty arises from the fact that it is not clear how to decide whether a given cross-free system is a skeleton or not. (Note that the size of \mathcal{O}^1 may be exponentially large.) To overcome this difficulty, we work with special kind of cross-free systems. Let us call a cross-free system $\mathcal{H} \subseteq \mathcal{O}^1$ **stable** if it fulfills the following property:

$$L \text{ crosses some element of } \mathcal{H} \text{ whenever } L \in \mathcal{O}^1 - \mathcal{H} \text{ and } \exists K \in \mathcal{H} : L \preceq K. \quad (2.1)$$

This means that if \mathcal{H} has an element larger than L , then L cannot be added to \mathcal{H} . Given a stable system, the following claim provides a straightforward way to decide whether it is a skeleton.

Claim 2.7. *A stable cross-free system is a skeleton if and only if it contains all the maximal members of \mathcal{O}^1 .*

Proof. On the one hand, any skeleton should contain all the maximal one-way pairs in \mathcal{O}^1 since a maximal one-way pair cannot cross any other set. On the other hand, for a contradiction, suppose that a stable system \mathcal{H} contains all the maximal members, yet it is not a skeleton. Choose an $L \notin \mathcal{H}$ with $\mathcal{H} \cup \{L\}$ cross-free. There is a maximal element $K \in \mathcal{O}^1$ with $L \preceq K$. By our assumption, $K \in \mathcal{H}$, contradicting the definition of stability. \square

Assume we are given a stable cross-free system \mathcal{H} which is not a skeleton. In the following, we investigate how a set $K \in \mathcal{O}^1 - \mathcal{H}$ can be found with the property that $\mathcal{H} \cup \{K\}$ is stable as well. As \mathcal{H} is not a skeleton, there is a maximal element M with $M \in \mathcal{O}^1 - \mathcal{H}$. Let

$$\mathcal{L}_1 := \{K \in \mathcal{H} : K \preceq M\}; \quad \mathcal{L}_2 := \{K \in \mathcal{H} : K \not\preceq M\} \quad (2.2)$$

We say that a one-way pair L **fits** the pair (\mathcal{H}, M) if (a) $L \in \mathcal{O}^1 - \mathcal{H}, L \preceq M$; (b) L is independent from all members of \mathcal{L}_2 and (c) either $K \preceq L$ or $K^- \cap L^- = \emptyset$ for every $K \in \mathcal{L}_1$.

Lemma 2.8. *If L is a minimal member of $\mathcal{O}^1 - \mathcal{H}$ fitting (\mathcal{H}, M) , then $\mathcal{H} + L$ is a stable cross-free system.*

This is a straightforward consequence of the following claim.

Claim 2.9. *Let $L \in \mathcal{O}^1 - \mathcal{H}, L \preceq M$. The following two properties are equivalent: (i) L fits (\mathcal{H}, M) ; (ii) $\mathcal{H} + L$ is cross-free.*

Proof. (i) \Rightarrow (ii) is straightforward. For the other direction we have to verify (b) and (c) of the above definition. By (2.1), either $K \preceq L$ or L and K are independent for every $K \in \mathcal{H}$. Assume now $K \preceq L$ for some $K \in \mathcal{L}_2$. In this case $K \preceq L \preceq M$, contradicting the definition of \mathcal{L}_2 . For (c) we need $K^- \cap L^- = \emptyset$ if K and L are independent for some $K \in \mathcal{L}_1$. This follows by $K, L \preceq M$, thus $K^+ \cap L^+ \supseteq M^+$. \square

Observe that M itself fits (\mathcal{H}, M) ensuring the existence of a one-way pair L satisfying the conditions of Lemma 2.8. So $K = L$ is an appropriate choice. Such an L can be found using bipartite matching theory. The description of this subroutine is quite technical and rather standard, therefore it is postponed to Section 2.4.

2.1.2 Description of the Dual Oracle

Given the above subroutine for constructing a skeleton, we have the following oracle to determine the value $\nu(D) = \nu(\mathcal{O}^1)$ in a $(k - 1)$ -connected digraph on n nodes: we construct a skeleton, then we apply Dilworth's theorem. (It is well-known that computing a maximum antichain

and a minimum chain-decomposition of a partially ordered set can be reduced to a maximum matching computation in a bipartite graph.) The size of the maximum antichain will give the value $\nu(D)$.

A trivial upper bound on the size of the optimal augmenting edge set – and by Theorem 1.5, also on the number of pairwise independent sets – is n^2 . A better bound can be given by Corollary 4.7 in [31]: there is an augmenting edge set consisting of pairwise node-disjoint circuits and paths, hence the optimum value is at most n . A chain can also have at most n elements, thus the cardinality of a skeleton is at most $s = n^2$.

As shown in the Section 2.4, if s is an upper bound on the size of a skeleton, then it can be constructed in time $O(n^5 + sn^4) = O(n^6)$. Finding a maximum antichain in a poset of size $O(s)$ can be reduced to finding a maximum matching in a bipartite graph on $O(s)$ nodes and $O(s^2)$ edges. Using the Hopcroft-Karp algorithm [69, Vol A., p. 264] this can be done in $O(s^{2.5})$ running time. This gives $O(n^5)$ for $s = n^2$, so the total running time of the Dual Oracle is $O(n^6)$.

As already indicated in the Introduction, the Dual Oracle may be used to compute the optimal augmentation. For this, we need to call the Dual Oracle at most n^2 times, thus the total complexity is $O(n^8)$. (For comparison, the running time of the algorithm in Chapter 4 is $O(n^7)$ for the same problem.)

However, the correctness of the present approach does rely on Theorem 1.5. In the next section we use a more direct approach for finding the optimal augmentation.

2.2 Algorithmic Proof of Theorem 1.6

In this section we give a proof of Theorem 1.6 and sketch another algorithm, which uses the Dual Oracle only once. After a skeleton \mathcal{K} is determined, an augmenting set of \mathcal{K} can be transformed to an augmenting set of the entire \mathcal{O}^1 . More precisely, we will prove the following:

Theorem 2.10. *For a crossing system \mathcal{F} and a one-way pair $K \in \mathcal{F}$, if an edge set F covers $\mathcal{F} \div K$, then there exists an F' covering \mathcal{F} with $|F'| = |F|$, and furthermore $\rho_{F'}(v) = \rho_F(v)$, $\delta_{F'}(v) = \delta_F(v)$ for every $v \in V$.*

We begin with the definition of the elementary augmenting step. Consider a crossing family $\mathcal{F} \subseteq \mathcal{O}^1$ and $F \subseteq V^2$. An edge $uv \in V^2 - F$ is **bad** (with respect to \mathcal{F} and F) if there exists an $L \in \mathcal{F}$ covered by uv , but not covered by F . Let $W(F) = W_{\mathcal{F}}(F)$ denote the set of bad edges.

Consider an augmenting edge set F of $\mathcal{F}' := \mathcal{F} \div K$. For two edges $x_1y_1, x_2y_2 \in F$, by **flipping** (x_1y_1, x_2y_2) , we mean replacing F by $F' = (F - \{x_1y_1, x_2y_2\}) \cup \{x_1y_2, x_2y_1\}$. A flipping is called **improving** if F' augments a strictly larger subset of \mathcal{F} than F does. Note that this is equivalent to requiring that $W(F') \subsetneq W(F)$. Since the total number of edges is n^2 , we obtain that after at most n^2 improving flippings the resulting subset of edges must augment

the whole \mathcal{F} . The following lemma, which is the heart of the proof of Theorem 1.6 and the algorithm, asserts the existence of an improving flipping.

Lemma 2.11. *Let $\mathcal{F} \subseteq \mathcal{O}^1$ be a crossing family. Let K be a member of \mathcal{F} and F an augmenting edge set of $\mathcal{F}' := \mathcal{F} \div K$. If F does not augment \mathcal{F} , then there is an improving flipping.*

Proof. Let us choose two (not necessarily distinct) members X and Y of \mathcal{F} that are not covered by F so that $X \preceq Y$, X is minimal (in the sense that X' is covered by F for every $X' \in \mathcal{F}, X' \prec X$), while Y is maximal in an analogous sense.

Since F does not cover X and Y , we have $X, Y \in \mathcal{F} - \mathcal{F}'$, that is, both X and Y cross K . Therefore $X \wedge K \prec X$ and $Y \vee K \succ Y$. By the minimality of X , $X \wedge K$ is covered by F , that is, there is an edge $x_1y_1 \in F$ covering $X \wedge K$. Since F does not cover X , we must have $x_1 \in X^- \cap K^-$ and $y_1 \in K^+ - X^+$. Analogously, there is an edge $x_2y_2 \in F$ covering $Y \vee K$ for which $x_2 \in K^- - Y^-$, $y_2 \in Y^+ \cap K^+$.

Let F' be the edge set resulting by flipping (x_1y_1, x_2y_2) . We are going to show that this flipping is improving. Since X is covered by F' but not covered by F , we only have to show that every member of \mathcal{F} covered by F is covered by F' , as well.

Suppose indirectly that there is a member M of \mathcal{F} which is covered by F but not by F' . In particular, no element of $F - \{x_1y_1, x_2y_2\}$ covers M . It is not possible that both x_1y_1 and x_2y_2 cover M since then both x_1y_2 and x_2y_1 would also cover M , that is, F' would cover M . Therefore there is exactly one element in F covering M and this only element is either x_1y_1 or x_2y_2 . Let us assume first that M is covered by x_1y_1 .

Claim 2.12. *Y and M are dependent.*

Proof. For a contradiction, suppose that Y and M are independent. $K \wedge Y$ and M are dependent as x_1y_1 covers both. Thus we can apply Lemma 2.5(i) with $L_1 = K, L_2 = Y, L_3 = M$ giving $K^- - Y^- \subseteq M^-$. This is contradiction since $x_2 \in K^- - Y^-$ and $x_2 \notin M^-$ as x_2y_1 does not cover M . \square

By the above claim we know that $Y \vee M \in \mathcal{F}$. The assumption that M is not covered by x_1y_2 gives $y_2 \in Y^+ - M^+$, thus $M \not\preceq Y$, implying $Y \cup M \succ Y$. By the maximality of Y , $Y \vee M$ is covered by an element xy of F and xy is different from both x_1y_2 and x_2y_1 since $y_1, y_2 \notin (M \vee Y)^+$. By Lemma 2.2, xy covers either M or Y . However, $xy \in F' \cap F$ and hence xy covers neither M nor Y , a contradiction.

The case when M is covered only by x_2y_2 also leads to contradiction by a similar argument using Lemma 2.5(ii). \square

Proof of Theorem 1.6. $\nu \leq \tau$ is straightforward. The proof of $\nu \geq \tau$ is by induction on $|\mathcal{F}|$. If \mathcal{F} is cross-free, applying Dilworth's theorem to the partially ordered set (\mathcal{F}, \subseteq) , we obtain that there is a maximum subfamily \mathcal{I} of \mathcal{F} consisting of pairwise incomparable members and that \mathcal{F}

can be decomposed into $\gamma := |\mathcal{I}|$ chains. Since \mathcal{F} is assumed to be cross-free, the members of \mathcal{I} are pairwise independent. Furthermore, it is easy to see that the chain-decomposition of \mathcal{F}_s corresponds to a set F of γ edges covering \mathcal{F} . Hence we obtained the required covering F of \mathcal{F} and independent subfamily \mathcal{I} of \mathcal{F} for which $|F| = |\mathcal{I}|$.

Assume now \mathcal{F} contains crossing one-way pairs K and K' . Let $\mathcal{F}' = \mathcal{F} \div K$, a crossing system by Lemma 2.4. As $K' \notin \mathcal{F}'$, we may apply the inductive statement for \mathcal{F}' giving an edge set F covering \mathcal{F}' among $|F|$ pairwise independent one-way pairs. The proof is finished using Lemma 3.11. \square

2.2.1 Description of the Algorithm

Our next goal is to transform the inductive proof above into an algorithm, that constructs an independent subset \mathcal{I} of \mathcal{O}^1 and an covering edge set F of \mathcal{O}^1 so that $|\mathcal{I}| = |F|$. It consists of two phases.

In **Phase 1** our algorithm uses the Dual Oracle. It determines a skeleton $\mathcal{K} = \{K_1, \dots, K_\ell\}$, and by Dilworth's theorem it finds a maximum antichain along with a minimum chain-decomposition. The chain-decomposition of \mathcal{K} corresponds to a subset F' of edges covering \mathcal{K} for which $|F'| = |\mathcal{I}|$. The antichain \mathcal{I} will be output by the whole algorithm as a maximum cardinality independent subset of \mathcal{O}^1 .

Phase 2 will terminate by outputting a covering of \mathcal{O}^1 of cardinality $|\mathcal{I}|$. Let $\mathcal{F}_0 = \mathcal{O}^1$ and $\mathcal{F}_j := \mathcal{O}^1 \div \{K_1, \dots, K_j\}$ for each $j = 1, \dots, \ell$. From Phase 1, we have $\mathcal{F}_\ell = \mathcal{K}$ covered. By Lemma 2.11, when applied to $\mathcal{F}_{\ell-1}, \mathcal{F}_\ell, K_\ell$ in place of $\mathcal{F}, \mathcal{F}', K$, respectively, we can find an improving flipping and obtain a revised covering F'' of \mathcal{F}_ℓ which covers a strictly larger subset of $\mathcal{F}_{\ell-1}$ as F' does. Since the number of bad edges is at most n^2 and an improving flipping reduces this number, after at most n^2 improving flippings the resulting covering of \mathcal{F}_ℓ will cover $\mathcal{F}_{\ell-1}$. Then we can iterate this step with $\mathcal{F}_{\ell-2}, \mathcal{F}_{\ell-1}, K_{\ell-1}, \dots, \mathcal{F}_0, \mathcal{F}_1, K_1$, and finally we get a cover F' of $\mathcal{O}^1 = \mathcal{F}_0$. F' will be the output of the algorithm as a minimal augmenting edge set of D .

We have outlined the steps of the algorithm and proved its validity. Phase 1 can be preformed as described in Section 2.1.1. For the realization of Phase 2, we can use similar techniques. However, we omit this analysis. Our reason for this is that the analysis is quite technical, and we could not improve on the running time bound of the Dual Algorithm.

2.3 Further remarks

2.3.1 Node-induced cost functions

The cost function $c : A \rightarrow \mathbb{R}$ is called **node-induced** if there exists two cost functions $c^-, c^+ : V \rightarrow \mathbb{R}$ so that $c(uv) = c^-(u) + c^+(v)$ for each edge $uv \in A$. Given a node-induced cost function c , a cover of a skeleton \mathcal{K} can be extended to a cover of \mathcal{O}^1 of the same cost by Theorem 2.10. Therefore the only task left is to determine a minimum cost cover of a skeleton.

Finding a minimum cardinality cover of a skeleton was an application of Dilworth's theorem. As already mentioned, this can be deduced to finding a maximum matching in a bipartite graph. Analogously, we show that finding a minimum cost cover (for node-induced costs) goes back to finding a maximum cost matching in a bipartite graph by using the standard reduction.

For the poset (\mathcal{K}, \preceq) , construct a bipartite graph $G = (A, B; E)$ so that to each element $K \in \mathcal{K}$ we have corresponding nodes $k' \in A$, $k'' \in B$, and if $K \preceq L$ then $k'l'' \in E$. Given a matching M , a chain cover of size $n - |M|$ can be obtained as follows. Starting from an uncovered node $k'_1 \in A$, if k''_1 is uncovered by M , then let the singleton chain $\{K_1\}$ correspond to k'_1 . Otherwise, let k'_2 be the node so that $k''_1 k'_2 \in M$, and define k_{i+1} so that if k''_i is covered by M , then $k''_i k'_{i+1} \in M$. This defines a chain $K_1 \preceq K_2 \preceq \dots \preceq K_\ell$, and these chains are pairwise disjoint if starting for different uncovered members of M .

Given the cost functions c^- and c^+ on V , define $w(k') = \min_{v \in K^-} c^-(v)$ and $w(k'') = \min_{v \in K^+} c^+(v)$. Observe that minimum cost of an edge covering the chain constructed above is exactly $w(k'_1) + w(k''_\ell)$. Therefore, if we consider the cost function on E induced by this w , then a matching M corresponds to a chain cover of cost equal to the total cost of the uncovered nodes. Hence finding a minimum cost chain cover is equivalent to finding the maximum cost of a matching, solvable via the Hungarian Method.

2.3.2 Generalization to Theorem 1.2

The proof of Theorem 1.6 given in Section 2.2 can also be extended to a new, algorithmic proof of the more general Theorem 1.2. Here we give only a brief sketch of this rather technical argument, detailed in [72, Section 4.4.2].

Unfortunately, Theorem 2.1 is not true in general for arbitrary crossing family \mathcal{F} in place of \mathcal{O}^1 . The main reason is that the innocent-looking Claim 2.3 fails to hold: there might exist set pairs $M \neq N$ with $M^- \subseteq N^-$, $M^+ \subseteq N^+$. Of course, in such a case one might argue that N is superfluous since if an edge set covers M , then it automatically covers N . Yet we cannot simply leave all such pairs N from \mathcal{F} as we may end up with a family of set pairs which is not crossing.

A possible solution is the following. Let us call a pair N **slim** if no other pair $M \in \mathcal{F}$ with $M^- = N^-$, $M^+ \subsetneq N^+$ exists. (It is still possible that there is an M with $M^- \subsetneq N^-$,

$M^+ = N^+.$)

We modify the definition of stability so that \mathcal{H} is stable if it is cross-free, each element of \mathcal{H} is slim, and instead of (2.1), it satisfies

$$L \text{ is either not slim, or crosses some element of } \mathcal{H} \\ \text{whenever } L \in \mathcal{O}^1 - \mathcal{H} \text{ and } \exists K \in \mathcal{H} : L \preceq K.$$

Given a stable $\mathcal{K} = \{K_1, \dots, K_\ell\}$ maximal for containment, so that each set $\{K_1, \dots, K_i\}$, $i = 1, \dots, \ell$ is stable, it can be proved that a cover of \mathcal{K} can be transformed to a cover of \mathcal{F} . ST -edge-connectivity augmentation by one can be tackled by this approach.

It would be highly desirable to extend these methods for Theorem 1.1, since it could give a simpler alternative to the currently existing only combinatorial algorithm for directed connectivity augmentation (the one in Chapter 4). Moreover, it could be possibly extendable to a polynomial time algorithm. (The algorithm in Chapter 4 is pseudopolynomial.) Unfortunately, we could not find such an extension so far: we do not even have a good idea how skeletons in \mathcal{S} should be defined.

2.4 Implementation via bipartite matching

In this section we present how the subroutine for constructing a skeleton can be implemented using bipartite matching theory. Given the $(k-1)$ -connected digraph $D = (V, A)$, let us construct the bipartite graph $B = (V', V''; H)$ as follows. With each node $v \in V$ associate nodes $v' \in V'$ and $v'' \in V''$ and an edge $v'v'' \in H$. With each edge $uv \in V$ associate an edge $u'v'' \in H$. For a set $X \subseteq V$, we denote by X' and X'' its images in V' and V'' , respectively. The $(k-1)$ -connectivity of G implies that B is $(k-1)$ -**elementary bipartite**, that is, for each $\emptyset \neq X' \subseteq V'$, either $\Gamma(X') = V''$ or $|\Gamma(X')| \geq |X'| + k - 1$. (See Section 1.1.2 on k -elementary bipartite graphs.) We say that $X' \subseteq V'$ is **tight** if $|\Gamma(X')| = |X'| + k - 1$ and $\Gamma(X') \neq V''$. Let \mathcal{R} denote the set of tight sets. Observe that $X' \in \mathcal{R}$ if and only if $X \in \mathcal{O}^1$. In this context, we say that an edge $x'y''$ **covers** the tight set X' if $x' \in X'$, $y'' \in V'' - \Gamma(X')$, or equivalently, if the edge xy covers the one-way pair X .

Given a function $f : V' \cup V'' \rightarrow \mathbb{N}$ we call the set $F \subseteq H$ an f -**factor** if $d_F(x) = f(x)$ for every $x \in V' \cup V''$. Let $f(Z) = \sum_{x \in Z} f(x)$ for $Z \subseteq V' \cup V''$.

Claim 2.13. *Consider a bipartite graph $B = (V', V''; H)$ and a function $f : V' \cup V'' \rightarrow \mathbb{N}$ so that $f(V') = f(V'')$ and $f(x) = 1$ or $f(y) = 1$ for every $xy \in H$. An f -factor exists if and only if $f(X) \leq f(\Gamma(X))$ for every $X \subseteq V'$.*

Proof. An easy consequence of Hall's theorem, replacing each $x \in V' \cup V''$ by $f(x)$ copies. The condition $f(x) = 1$ or $f(y) = 1$ for every $xy \in H$ guarantees that at most one copy of the same edge may be used. \square

First we show how the maximal elements of \mathcal{R} can be found; this in turn provides the maximal elements of \mathcal{O}^1 . Let us consider nodes $u' \in V'$, $v'' \in V''$ with $u'v'' \notin H$. A tight set $X' \in \mathcal{R}$ is called an $u\hat{v}$ -set if $u' \in X'$ and $v'' \notin \Gamma(X')$. For an edge $u'v'' \notin H$, consider the following f . Let $f(u') = f(v'') = k + 1$ and for $z \in (V' - u') \cup (V'' - v'')$, let $f(z) = 1$. An f -factor for this f is called a k - uv -factor. If B is a $(k - 1)$ -elementary bipartite graph, then Claim 2.13 implies the existence of a $(k - 1)$ - uv -factor. Let F_{uv} denote one of them.

Claim 2.14. *If there is a k - uv -factor, then there exists no $u\hat{v}$ -set.*

Proof. Assume X' is a $u\hat{v}$ -set. As $X' \in \mathcal{R}$, $|\Gamma(X')| = |X'| + k - 1$. Since $u' \in X'$, $v'' \notin \Gamma(X')$, we have $f(X') = |X'| + k$, $f(\Gamma(X')) = |X'| + k - 1$, thus no k - uv -factor may exist. \square

It is easy to see that any two $u\hat{v}$ -sets are dependent and the union and intersection of two $u\hat{v}$ -sets are $u\hat{v}$ -sets as well. Thus if the set of $u\hat{v}$ -sets is nonempty, then it contains unique minimal and maximal elements. In what follows we show how these can be found algorithmically. For an edge set $F \subseteq H$, we say that the path $U = x_0y_0x_1y_1 \dots x_t y_t$ is an **alternating path** for F from x_0 to y_t , if $x_i \in V'$, $y_i \in V''$, $x_i y_i \in H - F$ for $i = 0, \dots, t$, and $y_i x_{i+1} \in F$ for $i = 0, \dots, t - 1$. Under the same conditions we also say that $x_0 y_0 x_1 y_1 \dots x_t$ is an alternating path for F from x_0 to x_t .

Claim 2.15. (a) *If there exists an alternating path for F_{uv} from u' to v'' , then there exists no $u\hat{v}$ -set.* (b) *Assume there is no alternating path for F_{uv} from u' to v'' ; let S_1 denote the set of nodes $z \in V$ having an alternating path for F_{uv} from u' to z' . Then S_1' is the unique minimal $u\hat{v}$ -set.* (c) *Assume no alternating path exists for F_{uv} from u' to v'' ; let S_2 denote the set of nodes $z \in V$ having an alternating path for F_{uv} from z' to v'' . Then $V' - S_2'$ is the unique maximal $u\hat{v}$ -set.*

Proof. (a) Let U be an alternating path for F_{uv} from u' to v'' . Then $F \Delta U$ is a k - uv -factor so by Claim 2.14, no $u\hat{v}$ -set exists. (b) Let Z' be an arbitrary $u\hat{v}$ -set. For every $x' \in Z' - u'$, $\Gamma(Z')$ contains a unique y'' with $x'y'' \in F_{uv}$. The number of $y'' \in V''$ with $u'y'' \in F_{uv}$ is exactly k , and all of them are contained in $\Gamma(Z')$. These are $|Z'| + k - 1$ different elements of $\Gamma(Z')$, and since $Z' \in \mathcal{R}$, $\Gamma(Z')$ has no elements other than these. This easily implies that Z' contains every $x' \in V'$ for which there is an alternating path for F_{uv} from u' to x' , showing $S_1' \subseteq Z'$. It is left to prove that $S_1' \in \mathcal{R}$. From the definition of S_1' , it follows that for every $y'' \in \Gamma(S_1')$, there exists an $x' \in S_1'$ with $x'y'' \in F_{uv}$, proving $|\Gamma(S_1')| = |S_1'| + k - 1$. The proof of (c) follows the same lines. \square

For the initialization of the algorithm, we determine the edge sets F_{uv} by a single max-flow computation for every $u' \in V'$, $v'' \in V''$, $u'v'' \notin H$. By Claim 2.15, the maximal $u\hat{v}$ -sets can be found by a breadth-first search. The maximal ones among them correspond to the maximal

elements of \mathcal{O}^1 (note that the maximal $u\hat{v}$ -set might be contained in some other $x\hat{y}$ -set). We will use the sets F_{uv} also in the later steps of the algorithm.

Up to this point, all results will be applicable almost word for word for undirected augmentation in Section 3.5. The next part will also follow roughly the same lines, but there will be certain differences according to the different notion of stability in the two cases.

To implement the basic step of the algorithm, consider a stable cross-free system \mathcal{H} which is not a skeleton, a maximal element $M \in \mathcal{O}^1 - \mathcal{H}$ and $\mathcal{L}_1, \mathcal{L}_2$ as defined by (2.2). Our task is to find a K fitting (\mathcal{H}, M) and minimal subject to this property. Let \mathcal{T} be the set of the maximal elements of \mathcal{L}_1 .

Claim 2.16. *\mathcal{T} consists of pairwise tail-disjoint one-way pairs.*

Proof. Let $T_1, T_2 \in \mathcal{T}$. As they are maximal, they cannot be comparable, thus either $T_1^- \cap T_2^- = \emptyset$ or $T_1^+ \cap T_2^+ = \emptyset$. The latter is excluded since $T_1, T_2 \preceq M$ implies $T_1^+ \cap T_2^+ \supseteq M^+$. \square

Let us construct $B_1 = (V', V''; H_1)$ from B by adding some new edges as follows. For each $K \in \mathcal{L}_2$, add the edge $x'y'' \in H_1$ for every $x \in K^-, y \in K^+$. Furthermore, let $x'y'' \in H_1$ whenever $T \in \mathcal{T}, x \in T^-, y \in V'' - T^+$.

Claim 2.17. *Let $L \in \mathcal{O}^1 - \mathcal{H}, L \preceq M$. Then L fits (\mathcal{H}, M) if and only if L' is a tight set in B_1 .*

Proof. Clearly, L' is tight in B_1 if and only if $L' \in \mathcal{R}$ and there is no new edge $x'y'' \in H_1 - H$ with $x' \in L'$ and $y'' \in V'' - \Gamma(L')$.

L fits (\mathcal{H}, M) if it is independent from all elements of \mathcal{L}_2 , and for arbitrary $T \in \mathcal{T}$, either $T^- \cap L^- = \emptyset$ or $T^- \subsetneq L^-$. If it satisfies these properties, no new edge in $H_1 - H$ covers L' , thus L' is tight also in B_1 . For the other direction, if L is dependent from some $K \in \mathcal{L}_2$, then there exists $x \in K^- \cap L^-, y \in K^+ \cap L^+$ with $x'y'' \in H_1$ covering L' . If for some $T \in \mathcal{T}, T$ would cross L , then by Claim 2.3, $L^+ - T^+ \neq \emptyset$, thus there exist $x \in T^- \cap L^-, y \in L^+ - T^+$ with $x'y'' \in H_1$ covering L . \square

To find an L as in Lemma 2.8, we need to add some further edges to B_1 to ensure that $L \in \mathcal{O}^1 - \mathcal{H}$. (Note that the elements of \mathcal{T} are all tight in B_1 .) Let $Q \subseteq M^-$ be an arbitrary (not necessarily tight) set. Let $Z(Q)$ denote the unique minimal K satisfying the following property:

$$K \in \mathcal{O}^1, Q^- \subseteq K^-, \text{ and } K \text{ fits } (\mathcal{H}, M). \quad (2.3)$$

We will determine $Z(Q)$ for different sets Q in order to find an appropriate L . $Z(Q)$ is well-defined since M itself satisfies (2.3); and if K and K' satisfy (2.3), then K and K' are dependent and it is easy to see that $K \cap K'$ also satisfies (2.3). The next claim gives an easy algorithm for finding $Z(Q)$ for a given Z .

Claim 2.18. Fix some $u \in Q$, $v \in M^+$. Let B_2 denote the graph obtained from B_1 by adding all edges $u'y''$ with $y'' \in \Gamma(Q')$. Let S denote the set of nodes $z \in V$ for which there exists an alternating path for F_{uv} from u' to z' . Then $Z(Q) = S$.

Proof. As M' is an $u\hat{v}$ -set in B_2 , applying Claim 2.15(a) for B_2 instead of B , we get that B_2 contains no alternating path for F_{uv} from u' to v'' . By Claim 2.15(b), S' is the unique minimal $u\hat{v}$ -set in B_2 . The new edges in B_2 ensure that $\Gamma(S \cup Q) = \Gamma(S)$, thus $Q \subseteq S$ is an easy consequence of Claim 2.3. By Claim 2.17, S is the unique minimal set satisfying (2.3), thus $Z(Q) = S$. \square

Let W denote the union of the tails of the elements of \mathcal{T} . First, we shall find a one-way pair L_1 fitting (\mathcal{H}, M) and $L_1^- - W \neq \emptyset$. Let us compute the set $Z(\{u\})$ for any $u \in M^- - W$. By Claim 2.18, this can be done by a single breadth-first search. An arbitrary minimal element of the set $\{Z(\{u\}) : u \in M^- - W\}$ is an appropriate choice for L_1 .

Thus L_1 can be found by $|M^- - W| = O(n)$ breadth-first searches. Now either L_1 is itself a minimal set fitting (\mathcal{H}, M) , or there exists an L_2 with $L_2^- \subseteq W \cap L_1^-$, also fitting (\mathcal{H}, M) . This is impossible if $T \preceq L_1$ holds for at most one $T \in \mathcal{T}$, and thus L_1 is a minimal set fitting (\mathcal{H}, M) in this case.

Assume now $T \preceq L_1$ holds for at least two different $T \in \mathcal{T}$. In order to obtain L_2 , let us compute $Z(T_i^- \cup T_j^-)$ for any $T_i, T_j \in \mathcal{T}$, $T_i \neq T_j$, $T_i, T_j \prec L_1$. Choosing a minimal one among these gives a minimal L_2 fitting (\mathcal{H}, M) . This can be done by performing $O(n^2)$ breadth-first searches.

As L_2 fits (\mathcal{H}, M) and is minimal subject to this property, $L := L_2$ is an appropriate choice.

Complexity

In order to construct a skeleton, first we need n^2 Max Flow computations for the maximal members and the auxiliary graphs. The running time of adding a member to a stable cross-free system is dominated by $O(n^2)$ breadth first searches. Thus if s is an upper bound on the size of a skeleton, then we can find one in $O(n^5 + sn^4)$ time by using an $O(n^3)$ maximum flow algorithm and an $O(n^2)$ breadth first search algorithm. .

Chapter 3

Undirected node-connectivity augmentation

This chapter is devoted to the proof of Theorem 1.37. As indicated in the introduction, both the proof and the algorithm are closely related to those in Section 2 for directed connectivity augmentation. In Section 3.1, we define some basic concepts concerning relations of clumps and families of clumps. A main difference between the directed and undirected case is that the clumps admit no natural partial order. Still, we will introduce the notion of nestedness, an analogue of comparability. Two clumps are said to be crossing if they are neither independent nor nested. We will also be able to “uncross” such clumps, by referring to meets and joins of certain strict one-way pairs. Crossing and cross-free families and skeletons of clumps will correspond naturally to those of strict one-way pairs. A new type of difficulty is encountered due to large clumps. Fortunately, it turns out that large clumps are nested with every other clump they are dependent from.

Section 3.2 contains the proof of Theorem 1.37, using an argument analogous to the one in Section 2.2. The algorithm for constructing a skeleton is discussed in Section 3.3, resembling the one in Section 2.1.1. Finally, in Section 3.4 we solve the minimum cost version for node-induced cost functions, and discuss further possible generalizations and extensions as well.

3.1 Preliminaries

First we give a brief motivation of concepts related to clumps. In a $(k - 1)$ -connected graph G , we may have sets $B \subsetneq V$ with $|B| = k - 1$, so that $V - B$ has $t \geq 2$ connected components. The components of $V - B$ form a clump. Moreover, any partition of the components to at least two classes also forms a clump, since in the definition, the pieces are not required to be connected. In order to make G k -connected, we need to add at least $t - 1$ edges between different components of $V - B$. For $t = 2$, an arbitrary edge between the two components suffices, however the

situation is more complicated for $t \geq 3$. In this case, the set B is often called a **shredder** in the literature.

For a clump $X = (X_1, X_2, \dots, X_t)$, let $N_X = V - \bigcup_i X_i$. X is called **basic** if all pieces X_i are connected. The clump Y is **derived** from the basic clump X if each piece of Y is the union of some pieces of X . By $D(X)$ we mean the set of all clumps derived from X , while $D_2(X)$ is used for the set of small clumps derived from X . Let \mathcal{C} denote the set of all basic clumps. For a set $\mathcal{F} \subseteq \mathcal{C}$, $D(\mathcal{F})$ denotes the union of the sets $D(X)$ with $X \in \mathcal{F}$. The clumps being in the same $D(X)$ can easily be characterized (see e.g. [49, 50, 59]):

Claim 3.1. (i) Two clumps X and Y are derived from the same basic clump if and only if $N_X = N_Y$. (ii) If two basic clumps X and Y have a piece in common, then $X = Y$. \square

For a clump X and an edge set F , let F/X be the graph obtained from (V, F) by deleting N_X and shrinking the components X_i to single nodes. Let $c_F(X)$ denote the number of connected components of F/X . F **covers** X if F/X is connected, that is, $c_F(X) = 1$. To cover X , we need at least $|X| - 1$ edges of F between different components of X . If X is a small clump, then F covers X if and only if F connects X . We say that F covers (resp. connects) $\mathcal{H} \subseteq D(\mathcal{C})$ if it covers (resp. connects) all clumps in \mathcal{H} . Clearly, F is an augmenting edge set if and only if it covers $D(\mathcal{C})$. The following simple claim shows that in order to cover a set \mathcal{F} of clumps, it suffices to connect every small clump derived from the members of \mathcal{F} .

Claim 3.2. For an edge set $F \subseteq \binom{V}{2}$ and $\mathcal{F} \subseteq \mathcal{C}$, the following three statements are equivalent: (i) F covers \mathcal{F} ; (ii) F covers $D(\mathcal{F})$; and (iii) F connects $D_2(\mathcal{F})$. \square

We have already defined when two clumps are independent: if no edge in $\binom{V}{2}$ connects both. Two clumps are **dependent**, if they are not independent.

We say that two clumps $X = (X_1, \dots, X_t)$ and $Y = (Y_1, \dots, Y_h)$ are **nested** if $X = Y$ or there exist indices $1 \leq a \leq t$ and $1 \leq b \leq h$ so that $Y_i \subsetneq X_a$ for every $i \neq b$ and $X_j \subsetneq Y_b$ for every $j \neq a$. We call X_a the **dominant piece** of X with respect to Y , and Y_b the dominant piece of Y w.r.t X . The following important lemma shows that a large basic clump is automatically nested with any other basic clump (see also in [59]).

Lemma 3.3. Assume X is a large basic clump, and Y is an arbitrary basic clump. If X and Y are dependent then X and Y are nested.

To prove this, first we need two simple claims.

Claim 3.4. For the basic clumps $X = (X_1, \dots, X_t)$ and $Y = (Y_1, \dots, Y_h)$, $X_i \cap N_Y = \emptyset$ implies $X_i \subseteq Y_j$ for some $1 \leq j \leq h$. \square

Claim 3.5. Let $X = (X_1, \dots, X_t)$ and $Y = (Y_1, \dots, Y_h)$ be two different clumps both basic or both small. If $X_s \subsetneq Y_b$ for some $1 \leq s \leq t$, $1 \leq b \leq h$, then X and Y are nested with Y_b being the dominant piece of Y w.r.t X .




Figure 3.1: The nested clumps $X = (X_1, X_2, X_3)$ and $Y = (Y_1, Y_2, Y_3, Y_4)$ with dominant pieces X_1 and Y_1 .

Proof. Consider an $\ell \neq b$. $X_s \subseteq Y_b$ implies $d(X_s, Y_\ell) = 0$, thus $Y_\ell \cap N_X = \emptyset$. Hence $Y_\ell \subseteq X_a$ for some $a \neq s$ follows either by Claim 3.4 or by $t = 2$. We claim that this a is always the same independently from the choice of ℓ . Indeed, assume that for some $\ell' \notin \{b, \ell\}$, $Y_{\ell'} \subseteq X_{a'}$ with $a' \neq a$.

The same argument applied with changing the role of X and Y (by making use of $Y_\ell \subseteq X_a$) shows that $X_{a'} \subseteq Y_j$ for some j , giving $Y_{\ell'} \subseteq Y_j$, a contradiction. $X_i \subseteq Y_b$ for $i \neq a$ can be proved by changing the role of X and Y again. Thus X and Y are nested with dominant pieces X_a and Y_b . \square

Proof of Lemma 3.3. The dependence implies $X_1 \cap Y_1 \neq \emptyset$, $X_2 \cap Y_2 \neq \emptyset$ by possibly changing the indices. Let $x_i = |N_Y \cap X_i|$, $y_i = |N_X \cap Y_i|$, $n_0 = |N_X \cap N_Y|$. Then $k - 1 \leq |N(X_1 \cap Y_1)| \leq n_0 + x_1 + y_1$. Since $k - 1 = |N_Y| = n_0 + \sum_i y_i$ this implies $\sum_{i \neq 1} y_i \leq x_1$ and similarly $\sum_{i \neq 1} x_i \leq y_1$. The same argument for $X_2 \cap Y_2$ gives $\sum_{i \neq 2} y_i \leq x_2$ and $\sum_{i \neq 2} x_i \leq y_2$.

Thus we have $x_i = y_i = 0$ for $i \geq 3$. This gives $X_3 \cap N_Y = \emptyset$ and hence $X_3 \subseteq Y_i$ for some i by Claim 3.4. The nestedness of X and Y follows by the previous claim. \square

Beyond the close analogy between the argument of Chapter 2 and the present one, strict one-way pairs will also be directly applied. We will simply use “one-way pair” meaning strict one-way pair in the rest of this chapter. For each small clump $X = (X_1, X_2)$, the two corresponding one-way pairs (X_1, X_2) and (X_2, X_1) are called the **orientations** of X . By the orientations of a large clump X we mean all orientations of the small clumps in $D_2(X)$. For a one-way pair $K = (K^-, K^+)$, its reverse is $\overleftarrow{K} = (K^+, K^-)$, and \underline{K} denotes the corresponding small clump (note that $\underline{K} = \overleftarrow{\underline{K}}$).

The relation between covering in the directed and undirected sense is the following. If an undirected edge uv connects a small clump X , then the directed edge uv covers exactly one of its two orientations (in the directed sense).

Take two dependent small clumps $X = (X_1, X_2)$ and $Y = (Y_1, Y_2)$. We say that their orientations L_X and L_Y are **compatible** if they are dependent one-way pairs. Clearly, any two dependent one-way pairs admit compatible orientations, and if L_X and L_Y are compatible, then so are \overleftarrow{L}_X and \overleftarrow{L}_Y . X and Y are said to be **simply dependent** if for an orientation L_X of X , there is exactly one compatible orientation L_Y of Y , and **strongly dependent** if both possible choices of L_Y are compatible with L_X . (Note that the definition is independent of the choice of the orientation L_X). X and Y are strongly dependent if and only if $X_i \cap Y_j \neq \emptyset$ for every $i = 1, 2, j = 1, 2$. The following claim is easy to see.




Figure 3.2: Simply dependent one-way pairs (a), and strongly dependent ones (b).

Claim 3.6. *Two small clumps X and Y are nested if and only if for some orientations K_X and K_Y , $K_X \preceq K_Y$.* \square

We are ready to define uncrossing of basic clumps. By uncrossing the dependent one-way pairs K and L we mean replacing them by $K \wedge L$ and $K \vee L$ (which coincide with K and L if K and L are comparable). For dependent basic clumps X and Y , we define a set $\Upsilon(X, Y)$ consisting of two or four pairwise nested clumps in the analogous sense. If X and Y are nested, then let $\Upsilon(X, Y) = \{X, Y\}$. By Lemma 3.3, this is always the case if one of X and Y is large. For the small basic clumps X and Y , consider some compatible orientations L_X and L_Y . If X and Y are simply dependent then let $\Upsilon(X, Y) = \{\underline{L}_X \wedge L_Y, \underline{L}_X \vee L_Y\}$. (Although there are two possible choices for L_X and L_Y , the set $\Upsilon(X, Y)$ will be the same.) If they are strongly dependent, then L_X is also compatible \overleftarrow{L}_Y . In this case let $\Upsilon(X, Y) = \{\underline{L}_X \wedge L_Y, \underline{L}_X \vee L_Y, \underline{L}_X \wedge \overleftarrow{L}_Y, \underline{L}_X \vee \overleftarrow{L}_Y\}$. It is easy to see that the clumps in $\Upsilon(X, Y)$ are nested with X and Y and with each other in both cases. We will need the following submodular-type property, corresponding to Lemma 2.2:

Claim 3.7. *For dependent basic clumps X, Y , if an edge uv connects a clump in $\Upsilon(X, Y)$ then it connects at least one of X and Y .* \square

We say that two clumps are **crossing** if they are dependent but not nested. Again by Lemma 3.3, two basic clumps may be crossing only if both are small. A subset $\mathcal{F} \subseteq \mathcal{C}$ is called **crossing** if for any two dependent clumps $X, Y \in \mathcal{F}$, $\Upsilon(X, Y) \subseteq D(\mathcal{F})$. (The reason for assuming containment in $D(\mathcal{F})$ instead of \mathcal{F} is that $\Upsilon(X, Y)$ might contain non-basic clumps.) Note that \mathcal{C} itself is crossing. For a crossing system \mathcal{F} and a clump $K \in \mathcal{F}$, let $\mathcal{F} \div K$ denote the set of clumps in \mathcal{F} independent from or nested with K . Similarly, for a subset $\mathcal{K} \subseteq \mathcal{F}$, $\mathcal{F} \div \mathcal{K}$ denotes the set of clumps in \mathcal{F} not crossing any clump in \mathcal{K} . An $\mathcal{F} \subseteq \mathcal{C}$ is **cross-free** if it contains no crossing clumps, that is, any two dependent clumps in \mathcal{F} are nested. (Note that a cross-free system is crossing as well.) A cross-free \mathcal{K} is called a **skeleton** of \mathcal{F} if it is maximal cross-free in \mathcal{F} , that is, $\mathcal{F} \div \mathcal{K} = \mathcal{K}$. By Lemma 3.3, a skeleton of \mathcal{C} should contain every large clump. Let us now prove the counterpart of Lemma 2.4:

Lemma 3.8. *For a crossing system $\mathcal{F} \subseteq \mathcal{C}$ and $K \in \mathcal{F}$, $\mathcal{F} \div K$ is also a crossing system.*

Proof. Let $\mathcal{F}' = \mathcal{F} \div K$. If K is large then $\mathcal{F}' = \mathcal{F}$ by Lemma 3.3, therefore K is assumed being small in the sequel. Let us fix an orientation L_K of K . Take crossing basic clumps $X, Y \in \mathcal{F}'$. Again by Lemma 3.3, if a clump in $\Upsilon(X, Y)$ is not basic, then it is automatically in $D(\mathcal{F}')$. We consider all possible cases as follows.

(I) Both are nested with K . Choose orientations L_X and L_Y compatible with L_K (but not necessarily with each other). (a) If $L_X \preceq L_K \preceq L_Y$ or $L_Y \preceq L_K \preceq L_X$, then X and Y are nested by Claim 3.6. (b) Let $L_X, L_Y \preceq L_K$. If L_X and L_Y are dependent, then $L_X \wedge L_Y, L_X \vee L_Y \preceq L_K$. If L_X and \overleftarrow{L}_Y are dependent, then $L_X \wedge \overleftarrow{L}_Y \preceq L_K$ and $\overleftarrow{L}_K \preceq L_X \vee \overleftarrow{L}_Y$. These arguments show $\Upsilon(X, Y) \subseteq D(\mathcal{F}')$. (c) In the case of $L_X, L_Y \succeq L_K$, the claim follows analogously.

(II) Both X and Y are independent from K . By Claim 3.7, all clumps in $\Upsilon(X, Y)$ are independent from K .

(III) One of them, say X is nested with K , and the other, Y is independent from K . Let L_X be an orientation of X compatible with L_K and L_Y an orientation of Y compatible with L_X . By symmetry, we may assume $L_X \preceq L_K$. Now $L_X \wedge L_Y \preceq L_K$, and we show that $\overline{L_X \vee L_Y}$ is independent from K . L_Y being an arbitrary orientation compatible with L_X , these again imply $\Upsilon(X, Y) \subseteq D(\mathcal{F}')$. L_Y and L_K are independent, but $L_K^- \cap L_Y^- \neq \emptyset$, thus $L_K^+ \cap L_Y^+ = \emptyset$, hence the one-way pairs $L_X \vee L_Y$ and L_K are independent. We also need to show that $\overleftarrow{L_X \vee L_Y}$ and L_K are independent. Indeed, their dependence would imply $L_Y^+ \cap L_K^- \neq \emptyset$, $L_Y^- \cap L_K^+ \neq \emptyset$, contradicting the independence of K and Y . \square

Finally, the sequence K_1, K_2, \dots, K_ℓ of clumps is called a **chain** if they admit orientations L_1, L_2, \dots, L_ℓ with $L_1 \preceq L_2 \preceq \dots \preceq L_\ell$. If $u \in L_1^-, v \in L_\ell^+$ then the edge uv connects all members of the chain.

3.2 The proof of Theorem 1.37

For a crossing system $\mathcal{F} \subseteq \mathcal{C}$, let $\tau(\mathcal{F})$ denote the minimum cardinality of an edge set covering \mathcal{F} . Let $\nu(\mathcal{F})$ denote the maximum of $\text{def}(\Pi)$ over groves consisting of a shrub and bushes of clumps in $D(\mathcal{F})$. First, we give the proof of the following slight generalization of Theorem 1.37 based on two lemmas proved in the following subsections (cf. Theorem 1.6).

Theorem 3.9. *For a crossing system $\mathcal{F} \subseteq \mathcal{C}$, $\nu(\mathcal{F}) = \tau(\mathcal{F})$.*

The two lemmas are these:

Lemma 3.10. *For a cross-free system \mathcal{F} , $\nu(\mathcal{F}) = \tau(\mathcal{F})$.*

Lemma 3.11. *For a cross-free system \mathcal{F} , if an edge set F covers $\mathcal{F} \div K$, then there exists an F' covering \mathcal{F} with $|F'| = |F|$, and furthermore $d_{F'}(v) = d_F(v)$ for every $v \in V$.*

For the directed case in Chapter 2, the claim analogous to Lemma 3.10 was straightforward by Dilworth's theorem, while Lemma 3.11 is word-by-word the same as Theorem 2.10. Also, Theorem 3.9 derives from the lemmas the same way as Theorem 1.6.

The following theorem may be seen as a reformulation of this proof, however, it will be more convenient for the aim of the algorithm and to handle the minimum cost version for node induced cost functions.

Theorem 3.12. *For a crossing system $\mathcal{F} \subseteq \mathcal{C}$ and a skeleton \mathcal{K} of \mathcal{F} , $\nu(\mathcal{K}) = \nu(\mathcal{F})$. Furthermore, if an edge set F covers the skeleton \mathcal{K} of \mathcal{F} , then there exists an F' covering \mathcal{F} with $|F'| = |F|$ and $d_{F'}(v) = d_F(v)$ for every $v \in V$.*

Proof. Let $\mathcal{K} = \{K_1, \dots, K_\ell\}$. For $i = 1, \dots, \ell$, let $\mathcal{F}_i = \mathcal{F} \div \{K_1, \dots, K_i\}$. Lemma 3.8 implies that \mathcal{F}_i is a crossing system as well. $\mathcal{F}_\ell = \mathcal{K}$ since \mathcal{K} is a skeleton. By Lemma 3.10, \mathcal{K} admits a cover F_ℓ with $|F_\ell| = \tau(\mathcal{K}) = \nu(\mathcal{K})$. Applying Lemma 3.11 inductively for \mathcal{F}_{i-1}, K_i and F_i for $i = \ell, \ell - 1, \dots, 1$, we get a cover F_{i-1} of \mathcal{F}_{i-1} with $|F_{i-1}| = |F_i|$. Finally, F_0 is a cover of $\mathcal{F} = \mathcal{F}_0$, hence $\nu(\mathcal{F}) \leq |F_0| = |F_\ell| = \nu(\mathcal{K})$, implying the first part of the theorem. The identity of the degree sequences follows by the second part of Lemma 3.11. \square

3.2.1 Covering cross-free systems

This section is devoted to the proof of Lemma 3.10. The analogous statement in the case of directed connectivity augmentation simply follows by Dilworth's theorem, which is a well-known consequence of the König-Hall theorem on the size of a maximum matching in a bipartite graph. In contrast, Lemma 3.10 is deduced from Fleiner's theorem, which is proved via a reduction to the Berge-Tutte theorem on maximum matchings in general graphs.

We need the following notion to formulate Fleiner's theorem. A triple $P = (U, \preceq, M)$ is called a **symmetric poset** if (U, \preceq) is a finite poset and M a perfect matching on U with the

property that $u \preceq v$ and $uu', vv' \in M$ implies $u' \succeq v'$. The edges of M will be called **matches**. A subset $\{u_1v_1, \dots, u_kv_k\} \subseteq M$ is called a **symmetric chain** if $u_1 \preceq u_2 \preceq \dots \preceq u_k$ (and thus $v_1 \succeq v_2 \succeq \dots \succeq v_k$). The symmetric chains S_1, S_2, \dots, S_t **cover** P if $M = \bigcup S_i$.

A set $\mathcal{L} = \{L_1, L_2, \dots, L_\ell\}$ of disjoint subsets of M forms a **legal subpartition** if $uv \in L_i$, $u'v' \in L_j$, $u \preceq u'$ yields $i = j$, and no symmetric chain of length three is contained in any L_i . The value of \mathcal{L} is $val(\mathcal{L}) = \sum_i \left\lceil \frac{|L_i|}{2} \right\rceil$.

Theorem 3.13 (Fleiner, [20]). *Let $P = (U, \preceq, M)$ be a symmetric poset. The minimum number of symmetric chains covering P is equal to the maximum value of a legal subpartition of P .*

Note that the $\max \leq \min$ direction follows easily since a symmetric chain may contain at most two matches belonging to one class of a legal subpartition. This theorem gives a common generalization of Dilworth's theorem and of the well-known min-max formula on the minimum size edge cover of a graph (a theorem equivalent to the Berge-Tutte formula).

First we show that Lemma 3.10 is a straightforward consequence if \mathcal{F} contains only small clumps. Consider the cross-free family \mathcal{F} of clumps, and let U be the set of all orientations of one-way pairs in \mathcal{F} . The matches in M consist of the two orientations of the same clump, while \preceq is the usual partial order on one-way pairs. A symmetric chain corresponds to a chain of clumps. Since all clumps in a chain can be connected by a single edge, a symmetric chain cover gives a cover of \mathcal{F} of the same size. On the other hand, a legal subpartition yields a grove with a shrub and bushes consisting of the clumps corresponding to the one-way pairs in L_i .

Let us now turn to the general case when \mathcal{F} may contain large clumps as well. For an arbitrary set $A \subseteq V$, let $A^* = V - (A \cup N(A))$. An edge set F **semi-covers** the clump $X = (X_1, \dots, X_t)$ if F contains at least $|X| - 1$ edges connecting X , and furthermore each clump (X_i, X_i^*) is connected for $i = 1, \dots, t$. (Note that $X_i^* = \bigcup_{j \neq i} X_j$.) F semi-covers \mathcal{F} if it semi-covers every $X \in \mathcal{F}$. Although a semi-cover is not necessarily a cover, the following lemma shows that it can be transformed into a cover of the same size.

Lemma 3.14. *If F is a semi-cover of \mathcal{F} , then there exists an edge set H covering \mathcal{F} with $|F| = |H|$ and $d_H(v) = d_F(v)$ for every $v \in V$.*

Proof. We are done if F covers all clumps in \mathcal{F} . Otherwise, consider a clump $X \in \mathcal{F}$ semi-covered but not covered. X is large, since semi-covered small clumps are automatically covered. Since X is connected by at least $|X| - 1$ edges of F , there is an edge $e = x_1y_1 \in F$ connecting X with $c_F(X) = c_{F-e}(X)$. Each (X_i, X_i^*) is connected, hence we may consider an edge $x_2y_2 \in F$ connecting X with x_2y_2 being in a component of F/X different from the one containing x_1y_1 . Let $F' = F - \{x_1y_1, x_2y_2\} + \{x_1y_2, x_2y_1\}$ denote the flipping of x_1y_1 and x_2y_2 . Clearly, $c_{F'}(X) = c_F(X) - 1$. We show that $c_{F'}(Y) \leq c_F(Y)$ for every $Y \in \mathcal{F} - X$, hence by a sequence of such steps we finally arrive at an H covering \mathcal{F} .

Indeed, assume $c_{F'}(Y) > c_F(Y)$ for some $Y \in \mathcal{F}$. X and Y are dependent since at least one of x_1y_1 and x_2y_2 connects both. By Lemma 3.3, X and Y are nested; let X_a and Y_b denote their dominant pieces. The nodes x_1, y_1, x_2, y_2 lie in four different pieces of X and thus at least three of them are contained in Y_b . Consequently, $c_{F'}(Y) = c_F(Y)$ yields a contradiction. \square

In what follows, we show how a semi-cover F of \mathcal{F} can be found based on a reduction to Fleiner's theorem. For a basic clump $X = (X_1, \dots, X_t)$, let $u_i^X = (X_i, X_i^*)$, $v_i^X = (X_i^*, X_i)$ and $U^X = \{u_i^X, v_i^X : i = 1, \dots, t\}$. Let $U = \bigcup_{X \in \mathcal{F}} U^X$. We say that the members of U^X are of type X . Let the matching M consist of the matches $u_i^X v_i^X$; such a match is called an X -**match**.

If X is small ($t = 2$), then $u_1^X = v_2^X$ and $v_1^X = u_2^X$, thus $|U^X| = 2$. If X is large, then $|U^X| = 2t$. In this case, let u_1^X and v_1^X be called the **special one-way pairs** w.r.t X . $u_1^X v_1^X$ is called a **special match**. Note that it matters here, which piece of X is denoted by X_1 (arbitrarily chosen though). Let the partial order \preceq' on U be defined as follows. If x and y are one-way pairs of different type, then let $x \preceq' y$ if and only if $x \preceq y$ for the standard partial order \preceq on one-way pairs. If x and y are both of type X for a large clump X , then let $x \preceq y$ if either $x = u_1^X, y = v_i^X$, or $x = u_i^X, y = v_1^X$ for some $i > 1$. In other words, \preceq' is the same as \preceq except that x and y are incomparable whenever x and y are of the same type X , and neither of them is special.

Claim 3.15. $P = (U, \preceq', M)$ is a symmetric poset.

Proof. The only nontrivial property to verify is the transitivity of \preceq' : $x \preceq' y$ and $y \preceq' z$ implies $x \preceq' z$. This follows by the transitivity of \preceq unless x and z are different one-way pairs of the same type X , and neither of them is special. Thus X is a large clump and by possibly changing the indices, assume $x = u_2^X, z = v_3^X$. y could be of type X only if it were special, excluded by $x = u_2^X \not\preceq u_1^X$ and $z = v_3^X \not\preceq v_1^X$. Hence y is of a different type Y .

Assume first $y = u_i^Y$ for some i . Now $X_2 \subseteq Y_i \subseteq X_3^*$ thus $N_X \cap Y_i = \emptyset$, giving by Claim 3.4 $Y_i \subseteq X_j$ for some $j \neq 3$. Consequently, $X_2 = Y_i$, a contradiction as it would lead to $X = Y$ by Claim 3.1. Next, assume $y = v_i^Y$. $X_3 \subseteq Y_i \subseteq X_2^*$ gives a contradiction the same way. \square

The following simple claim establishes the connection between dependency of clumps and comparability in P .

Claim 3.16. In a cross-free system \mathcal{F} , the clumps $X, Y \in \mathcal{F}$ are dependent if and only if for arbitrary i, j , u_i^X is comparable with either u_j^Y or v_j^Y . \square

Consider a symmetric chain cover S_1, \dots, S_t and a legal subpartition $\mathcal{L} = \{L_1, L_2, \dots, L_\ell\}$ with $\text{val}(\mathcal{L}) = t$. Let us choose \mathcal{L} so that ℓ is maximal, and subject to this, $\bigcup_{i=1}^\ell L_i$ contains the maximum number special matches. A symmetric chain S_i naturally corresponds to a chain of the clumps (X_j, X_j^*) for $u_j^X v_j^X \in S_i$. These can be covered by a single edge; hence a symmetric chain cover corresponds to an edge set F of the same size. A symmetric chain may contain both

$u_j^X v_j^X$ and $u_{j'}^X v_{j'}^X$ for $j \neq j'$ only if $j = 1$ or $j' = 1$. Consequently, F is a semi-cover as there are at least $|X| - 1$ different edges in F connecting X , and all (X_j, X_j^*) 's are connected.

It is left to show that \mathcal{L} can be transformed to a grove Π with $\text{def}(\Pi) = \text{val}(\mathcal{L})$. For a clump X , let $B(X)$ denote the set of indices j with $u_j^X v_j^X \in \bigcup_i L_i$. Most efforts are needed to ensure that the bushes consist of small clumps; allowing large clumps would enable a simpler argument.

Claim 3.17. *For any clump X , the X -matches corresponding to $B(X)$ are either all contained in the same L_i or are all singleton L_i 's. $1 \in B(X)$ always gives the first alternative.*

Proof. There is nothing to prove for $|X| = 2$, so let us assume $|X| \geq 3$. As \mathcal{L} is chosen with ℓ maximal, if $u_j^X v_j^X \in L_i$ with $|L_i| > 1$, then there is an $u_h^Y v_h^Y \in L_i$ with u_h^Y comparable with either u_j^X or v_j^X . If $Y \neq X$, then Claim 3.16 gives that u_h^Y is also comparable with $u_{j'}^X$ or $v_{j'}^X$ for any $j' \in B(X)$. If $Y = X$ then either $j = 1$ or $h = 1$ follows, implying $u_{j'} v_{j'} \in L_i$ for every $j' \in B(X)$. This argument also shows that $1 \in B(X)$ leads to the first alternative. \square

Let $\beta(X) = i$ in the first alternative if L_i is not a singleton, and $\beta(X) = 0$ in the second alternative. Let \mathcal{I} denote the set of indices for which L_i is a singleton. Take a clump X with $\beta(X) = i > 0$ (and thus $i \notin \mathcal{I}$). Let us say that a piece X_j is a **dominant piece** of X , if for some $Y \neq X$ with $\beta(Y) = i$, X_j is the dominant piece of X w.r.t. Y . Let $U(X)$ denote the set of the indices of the dominant pieces of X ; note that the set $U(X) - B(X)$ is possibly nonempty.

Claim 3.18. *If $\beta(X) = i > 0$, then $|B(X)| \geq 2$ implies $|B(X) \cap U(X)| = \emptyset$.*

Proof. First assume $B(X) \cap U(X) \neq \emptyset$ and $|U(X)| \geq 2$. Consider a $j \in B(X) \cap U(X)$ and a $j' \in U(X) - \{j\}$, say, X_j is the dominant piece of X w.r.t. Y and $X_{j'}$ the one w.r.t. Y' with $\beta(Y) = \beta(Y') = i$. It is easy to see that L_i contains a symmetric chain of length three consisting of a Y -match, $u_j^X v_j^X$ and a Y' -match.

Thus $B(X) \cap U(X) \neq \emptyset$ implies $|U(X)| = 1$. Let $U(X) = \{j\}$. Assume again that X_j is the dominant piece of X w.r.t. Y with $\beta(Y) = i$. We claim that $1 \notin B(X)$. Indeed, if $1 \in B(X)$ and $j \neq 1$, then a Y -match, $u_j^X v_j^X$ and $v_1^X u_1^X$ would form a symmetric chain in L_i . If $j = 1$, then a Y -match, $u_1^X v_1^X$ and $v_h^X u_h^X$ forms a symmetric chain for arbitrary $h \in B(X) - \{1\}$.

Let us replace L_i by $L'_i = L_i - \{u_j^X v_j^X\} + \{u_1^X v_1^X\}$. By Claim 3.16, any element of L'_i is incomparable to any element of L_h for $h \neq i$. It is easy to verify that L'_i does not contain any symmetric chain of length three given that L_i did not contain any. This is a contradiction as \mathcal{L} was chosen containing the maximal possible number of special matches. \square

Let us construct the grove Π as follows. For any X with $\beta(X) = 0$, $B(X) \neq \emptyset$, let $\tilde{X} \in D(X)$ denote the clump consisting of pieces X_i with $i \in B(X)$ and the piece $\bigcup_{j \notin B(X)} X_j$. The latter set is nonempty since $1 \notin B(X)$ by Claim 3.17, thus $|\tilde{X}| - 1 = |B(X)|$. Define the shrub as $\mathcal{B}_0 = \{\tilde{X} : \beta(X) = 0\}$. For $i \notin \mathcal{I}$, let $\mathcal{B}_i = \{(X_j, X_j^*) : u_j^X v_j^X \in L_i\}$. The following easy claim completes the proof.

Claim 3.19. Π is a grove with $\text{def}(\mathcal{B}_0) = |\mathcal{I}|$ and $\text{def}(\mathcal{B}_i) = \left\lceil \frac{|L_i|}{2} \right\rceil$ if $i \notin \mathcal{I}$.

Proof. Since the elements of different L_i 's are pairwise incomparable, Claim 3.16 implies that clumps in different bushes are independent from each other and from those in \mathcal{B}_0 . Assume an edge $uv \in \binom{V}{2}$ covers three clumps in some \mathcal{B}_i . If these three clumps were derived from different basic clumps, then L_i would contain a symmetric chain of length three. Thus we need to have two clumps derived from the same basic clump X : uv covers (X_j, X_j^*) , $(X_{j'}, X_{j'}^*)$ and (Y_h, Y_h^*) for $\beta(X) = \beta(Y) = i$. This is also impossible since either X_j or $X_{j'}$ would need to be the dominant piece of X w.r.t Y , a contradiction to Claim 3.18. \square

3.2.2 The proof of Lemma 3.11.

First we need the following lemmas.

Lemma 3.20. *Assume that for three small clumps $X = (X_1, X_2)$, $Y = (Y_1, Y_2)$, $Z = (Z_1, Z_2)$, all four sets $X_1 \cap Y_1 \cap Z_1$, $X_1 \cap Y_2 \cap Z_2$, $X_2 \cap Y_1 \cap Z_2$, $X_2 \cap Y_2 \cap Z_1$ are nonempty. Then all of X , Y and Z are derived from the same basic clump (and thus none of them is basic itself).*

Proof. Let $X_c = N_X$, $Y_c = N_Y$, $Z_c = N_Z$. By A_s for a sequence s of three literals each 1,2 or c , we mean the intersection of the corresponding sets. For example, $A_{12c} = X_1 \cap Y_2 \cap Z_c$.

The conditions mean that the sets A_{111} , A_{122} , A_{212} , A_{221} are nonempty. $V - (A_{111} \cup N(A_{111})) \neq \emptyset$ as there is no edge between A_{111} and X_2 , thus $|N(A_{111})| \geq k - 1$ as G is $(k - 1)$ -connected. This implies

$$k - 1 \leq |A_{c11} \cup A_{1c1} \cup A_{11c} \cup A_{1cc} \cup A_{c1c} \cup A_{cc1} \cup A_{ccc}| \quad (3.1)$$

as $N(A_{111})$ is a subset of the set on the RHS. Let us take the sum of these types of inequalities for all A_{111} , A_{122} , A_{212} , A_{221} . This gives $4(k - 1) \leq S_1 + 2S_2 + 4|A_{ccc}|$, where S_1 is the sum of the cardinalities of the sets having exactly one c in their indices, while S_2 is the same for two c 's.

On the other hand, $|X_c| = |Y_c| = |Z_c| = k - 1$. This gives $3(k - 1) = S_1 + 2S_2 + 3|A_{ccc}|$. These together imply $S_1 = S_2 = 0$, $|A_{ccc}| = k - 1$. We are done by Claim 3.1 since $N_X = N_Y = N_Z = A_{ccc}$. \square

Proof of Lemma 3.11. Let $\mathcal{F}' = \mathcal{F} \div K$. If K is large then $\mathcal{F}' = \mathcal{F}$ by Lemma 3.3, therefore K will be assumed to be small with an orientation L_K .

If F covers \mathcal{F}' but not \mathcal{F} , then by Claim 3.2 there exists a small clump $X \in D_2(\mathcal{F}) - D_2(\mathcal{F}')$ not connected by F , thus X and K are crossing. Choose X with the orientation L_X compatible with L_K so that L_X is minimal to these properties w.r.t. \preceq (that is, there exists no other uncovered X' with orientation $L_{X'}$ compatible with L_K so that $L_{X'} \prec L_X$.) Choose Y not connected by F with $L_X \preceq L_Y$, and L_Y maximal in the analogous sense ($X = Y$ is allowed).

$L_X \wedge L_K$ and $L_Y \vee L_K$ are nested with L_K and thus connected by edges $x_1y_1, x_2y_2 \in F$ with $x_1 \in L_X^- \cap L_K^-, y_2 \in L_Y^+ \cap L_K^+$. As X and Y are not connected, $y_1 \in L_K^+ - L_X^+, x_2 \in L_K^- - L_Y^-$ follows. Let $F' = F - \{x_1y_1, x_2y_2\} + \{x_1y_2, x_2y_1\}$ denote the flipping of x_1y_1 and x_2y_2 . F' connects X and Y , and we shall prove that F' connects all small clumps in $D_2(\mathcal{F})$ connected by F . Hence after a finite number of such operations all small clumps in $D_2(\mathcal{F})$ will be connected, so by Claim 3.2, \mathcal{F} will be covered.

For a contradiction, assume there is a small clump S connected by F but not by F' . (S is not necessarily basic.) No edge in $F \cap F'$ may connect S , hence either exactly one of x_1y_1 and x_2y_2 connect it, or if both then x_1 and y_2 are in the same piece and y_1 and x_2 in the other piece of S . In this latter case, K and S are strongly dependent.

(I) First, assume that only x_1y_1 connects S , and choose the orientation L_S with $x_1 \in L_S^-, y_1 \in L_S^+$. We claim that L_S and L_Y are also dependent. Indeed, if they are independent, then Lemma 2.5(i) is applicable for $L_1 = L_K, L_2 = L_Y, L_3 = L_S$, since $L_K \wedge L_Y$ and L_S are dependent because x_1y_1 connects both. This gives $x_2 \in L_K^- - L_Y^- \subseteq L_S^-,$ that is, x_2y_1 connects S , a contradiction.

Hence we may consider the one-way pair $L_S \vee L_Y$. $L_S \vee L_Y$ is strictly larger than L_Y , as if $L_S \preceq L_Y$ held, then S would be connected by x_1y_2 . By the maximal choice of L_Y , $L_S \vee L_Y$ is connected by some edge $f \in F$. By Claim 3.7, f also connects S or Y , implying $f = x_1y_1$. This is a contradiction as $x_1 \in L_S^- \cup L_Y^-$ and $y_1 \notin L_S^+ \cap L_Y^+$.

(II) If x_2y_2 is the only edge connecting S , we may use the same argument by exchanging \vee and \wedge , X and Y , “minimal” and “maximal” everywhere and applying Lemma 2.5(ii) instead of (i).

(III) Finally, if both x_1y_1 and x_2y_2 cover S , let L_S be chosen with $x_1, y_2 \in L_S^-, y_1, x_2 \in L_S^+$. The argument in (I) may be applied with the only difference that at the end $f = x_2y_2$ is also possible. This gives $x_2 \in L_Y^+ \cap L_S^+,$ thus $x_2 \in L_X^+.$ Analogously, the argument in (II) applies for $\overleftarrow{L}_S,$ and we get $y_1 \in L_X^- \cap L_S^+,$ thus $y_1 \in L_X^-.$

Now the clumps K, S and X satisfy the condition in Lemma 3.20, witnessed by nodes $x_1, x_2, y_2, y_1.$ This contradicts the assumption that K was a basic clump. \square

3.3 The Algorithm

As outlined in Section 1.5, the algorithm will be a simple iterative application of a subroutine determining the dual optimum $\nu(G).$ Theorem 3.12 shows that $\nu(G) = \nu(\mathcal{K})$ for an arbitrary skeleton $\mathcal{K}.$ Given a skeleton $\mathcal{K}, \nu(\mathcal{K})$ can be determined based on Fleiner’s theorem: Theorem 3.13 admits a (linear time) reduction to maximum matching in general graphs, as described in Section 3.3.2. As in Chapter 2, the naive greedy approach fails due to the possibly exponential size of $\mathcal{C}.$ The solution will be again the notion of stability, however, significantly more complicated than in Section 2.1.1.

3.3.1 Constructing a skeleton

Let us first introduce some new notation concerning pieces. If the set $B \subseteq V$ is a piece of the basic clump X , then let B^\sharp denote X . Let \mathcal{Q} be the set of all (connected) pieces of all basic clumps, whereas \mathcal{Q}_1 the set of all (not necessarily connected) pieces of all clumps. For a subset $\mathcal{A} \subseteq \mathcal{Q}$, \mathcal{A}^\sharp is the set of corresponding basic clumps (e.g. $\mathcal{Q}^\sharp = \mathcal{C}$).

As for the directed case, now we define stability. A cross-free set of $\mathcal{H} \subseteq \mathcal{C}$ is **stable** if it fulfills the following:

U crosses some element of \mathcal{H} whenever $U \in \mathcal{C} - \mathcal{H}$ and $\exists K, K' \in \mathcal{H} : K, U, K'$ forms a chain.

The following simple claim will be used for handling chains of length three.

Claim 3.21. *For pieces $B_1, B_2, B_3 \in \mathcal{Q}_1$, if (i) $B_1 \subseteq B_2 \subseteq B_3$ or (ii) $B_1 \subseteq B_2$ and $B_3 \subseteq B_2^*$, then the corresponding clumps $B_1^\sharp, B_2^\sharp, B_3^\sharp$ form a chain. \square*

Clearly, $\mathcal{H} = \emptyset$ is stable, and every skeleton is stable as well. Let $\mathcal{M} \subseteq \mathcal{Q}$ denote the set of the pieces minimal for inclusion. Based on the following claim (an analogue of Claim 2.7), we will be able to determine when a stable cross-free system is a skeleton. The subroutine for finding the elements of \mathcal{M} will be given in Section 3.5 among other technical details of the algorithm.

Claim 3.22. *The stable cross-free system $\mathcal{H} \subseteq \mathcal{C}$ is a skeleton if and only if $\mathcal{M}^\sharp \subseteq \mathcal{H}$.*

Proof. On the one hand, every skeleton should contain \mathcal{M}^\sharp . Indeed, consider an $M \in \mathcal{M}$. M^\sharp cannot cross any $X \in \mathcal{C}$, as $\Upsilon(X, M^\sharp)$ would contain a clump with a piece being a proper subset of M .

On the other hand, assume \mathcal{H} is not a skeleton even though $\mathcal{M}^\sharp \subseteq \mathcal{H}$. Hence there exists a clump $U = (U_1, \dots, U_t) \in \mathcal{C} - \mathcal{H}$, not crossing any element of \mathcal{H} . Consider minimal pieces $M_1 \subseteq U_1$, $M_2 \subseteq U_2$. Then $M_1^\sharp, U, M_2^\sharp$ forms a chain by Claim 3.21(ii), contradicting the stability. \square

Assume \mathcal{H} is a stable cross-free system, but not a skeleton. In the following, we show how \mathcal{H} can be extended to a stable cross-free system larger by one. By the above claim, there is an $M \in \mathcal{M}$ with $M^\sharp \in \mathcal{C} - \mathcal{H}$. Let

$$\mathcal{L}_1 := \{X \in \mathcal{H} : X \text{ and } M^\sharp \text{ are nested}\}, \quad \mathcal{L}_2 := \{X \in \mathcal{H} : X \text{ and } M^\sharp \text{ are independent}\} \quad (3.2)$$

Claim 3.23. *If $\mathcal{L}_1 = \emptyset$, then $\mathcal{H} + M^\sharp$ is a stable cross-free system.*

Proof. Indeed, assume that for some $U \in \mathcal{C} - \mathcal{H}$ and $K \in \mathcal{H}$, $\mathcal{H} + U$ is cross-free, although K, U, M^\sharp forms a chain. Now K and M are dependent and thus nested, a contradiction. \square

In the sequel we assume $\mathcal{L}_1 \neq \emptyset$. The key concept of the algorithm will be “fitting”: as in the directed case, we shall define when a piece $Z \in \mathcal{Q}$ **fits** the pair (\mathcal{H}, M) . However, the definition is significantly more complicated, therefore we formulate the main lemma in advance (cf. Lemma 2.8):

Lemma 3.24. *Let C be a minimal member of $\mathcal{Q} - \bigcup \mathcal{H}$ fitting (\mathcal{H}, M) . Then $\mathcal{H} + C^\sharp$ is a stable cross-free system.*

There exists a C satisfying the conditions of this lemma, as according to the definition, the pieces of M^\sharp different from M (that is, the connected components of M^*) fit (\mathcal{H}, M) . Such a C can be found using standard bipartite matching theory similarly as in Chapter 2; the technical details are postponed to Section 3.5.

The minimality of M implies that for any $X \in \mathcal{L}_1$, the dominant piece of M^\sharp w.r.t. X is a connected component of M^* . One simple notion before giving the definition of fitting is the following. For pieces $B, C \in \mathcal{Q}$, we say that B **supports** C if $B \subseteq C \subseteq M^*$. $B \in \mathcal{Q}$ supports $Y \in \mathcal{C}$ if B supports some piece of Y ; $X \in \mathcal{C}$ supports $B \in \mathcal{Q}$ if a piece of X supports B .

Definition 3.25. The piece $C \in \mathcal{Q}$ **fits** the pair (\mathcal{H}, M) if

- (a) $C^\sharp \in \mathcal{C} - \mathcal{H}, C \subseteq M^*$.
- (b) There exists a $W \in \mathcal{L}_1$ supporting C .
- (c) Consider a clump $X \in \mathcal{L}_1$ with dominant piece X_a w. r. t. M^\sharp , and another piece X_i with $i \neq a$. Then either $X_i \subsetneq C$ or $X_i \cap C = \emptyset$, and if $X_a \cap C \neq \emptyset$ then $X_i \cap C^* = \emptyset$.
- (d) C^\sharp is independent from every $X \in \mathcal{L}_2$.

The proof of Lemma 3.24 is based on the following claim:

Claim 3.26. *Let $C \in \mathcal{Q} - \bigcup \mathcal{H}, C \subseteq M^*$ supported by some $W \in \mathcal{L}_1$. The following two properties are equivalent: (i) C fits (\mathcal{H}, M) ; (ii) $\mathcal{H} + C^\sharp$ is cross-free.*

Proof. First we show that (i) implies (ii). C^\sharp is independent from all pairs in \mathcal{L}_2 . Consider an $X \in \mathcal{L}_1$. C^\sharp and X cannot cross by Lemma 3.3 whenever X or C^\sharp is large, thus let us assume they both are small basic clumps, $X = (X_1, X_2)$ with X_2 being the dominant piece of X w.r.t. M^\sharp . If X and C^\sharp are dependent, then $X_1 \cap C \neq \emptyset$ or $X_2 \cap C \neq \emptyset$. In the first case, (c) implies $X_1 \subsetneq C$ hence nestedness follows by Claim 3.5. So let us assume $X_1 \cap C = \emptyset$. By the dependency, $X_1 \cap C^* \neq \emptyset$, contradicting $X_2 \cap C \neq \emptyset$ by the second part of (c).

Next, we show that (ii) implies (i). (a) and (b) are included among the conditions. For (c), consider an $X \in \mathcal{L}_1$ with dominant piece X_a w.r.t. M and another piece $X_i, i \neq a$. Notice that $X_i \subseteq M^*$. If X and C^\sharp are independent, then $X_i \cap C = \emptyset$ as otherwise an edge between $X_i \cap C$ and M would connect both. If they are dependent so that the dominant side of X w.r.t. C^\sharp is

different from X_i , then $X_i \subsetneq C$ or $X_i \cap C = \emptyset$ follows. Finally, if the dominant side is X_i , then C cannot be the dominant side of C^\sharp w.r.t. X (as it would imply $M \subseteq X_a \subseteq C$), thus $C \subsetneq X_i$. Now W, C^\sharp, X forms a chain by Claim 3.21(i), a contradiction to the stability of \mathcal{H} .

Assume next $X_a \cap C \neq \emptyset$ and $X_i \cap C^* \neq \emptyset$. X and C^\sharp are again dependent and thus nested, and as above, the dominant side of X cannot be X_i . C cannot be the dominant side of C^\sharp as $X_i \subseteq C$ would contradict $X_i \cap C^* \neq \emptyset$. Hence $C \subseteq X_i^*$. We get a contradiction again because of the chain W, C^\sharp, X .

Finally for (d), assume C^\sharp and $X \in \mathcal{L}_2$ are dependent. C cannot be the dominant piece of C^\sharp w.r.t. X as it would yield $X \in \mathcal{L}_1$. Consequently, $X_i \subseteq C^*$ for a non-dominant piece X_i of X w.r.t. C^\sharp , and thus by Claim 3.21(ii), W, C^\sharp, X forms a chain, a contradiction to stability. \square

Proof of Lemma 3.24. Using Claim 3.26, it is left to show that no chain C^\sharp, U, K may exist with $K \in \mathcal{H}$, $U \in \mathcal{C} - (\mathcal{H} + C^\sharp)$ so that $\mathcal{H} + C^\sharp + U$ is cross-free. Indeed, if such a chain existed, then C^\sharp and K would be dependent and thus nested. Let C' be the dominant piece of C^\sharp w.r.t. K . If $C' \neq C$ then by Claim 3.21(ii), W, C^\sharp, K is a chain, contradicting the stability of \mathcal{H} . (W is the clump supporting C ensured by (b).)

If $C' = C$, then for some pieces U_1 of U and K_1 of K , $K_1 \subsetneq U_1 \subsetneq C$. Now $U_1 \in \mathcal{Q} - \bigcup \mathcal{H}$, $U_1 \subseteq M^*$ and K supports U_1 . By making use of Claim 3.26, U_1 fits (\mathcal{H}, M) , a contradiction to the minimal choice of C . \square

3.3.2 Description of the Dual Oracle

To determine the value of $\nu(G)$, we first construct a skeleton \mathcal{K} as described above. For \mathcal{K} , we apply the reduction to Theorem 3.13 as in Section 3.2.1. As already mentioned, a minimal chain decomposition along with maximal legal subpartition of a symmetric poset $P = (U, \preceq, M)$ may be found via a reduction to finding a maximum matching. For the sake of completeness and also because it will be needed for the minimum cost version, we include this reduction. Define the graph $C = (U, H)$ with $uv' \in H$ if and only if $u \prec v$ and $vv' \in M$ for some $v \in U$.

It is easy to see that the set $\{m_1, m_2, \dots, m_\ell\} \subseteq M$ is a symmetric chain if and only if there exists edges $e_1, \dots, e_{\ell-1} \in H$ such that $m_1 e_1 m_2 e_2 \dots m_{k-1} e_{k-1} m_k$ is a path, called an M -alternating path. The transitivity of \preceq ensures that $M \cup H$ contains no M -alternating cycles. Let $N \subseteq H$ be a matching in C . Then the components of $M \cup N$ are M -alternating paths, each containing exactly two nodes not covered by N . Hence finding a maximum matching in H is equivalent to finding a minimum chain cover in P . The running time of the most efficient maximum matching algorithm for a graph on n_1 nodes with m_1 edges is $O(\sqrt{n_1} m_1)$ [69, Vol I, p. 423].

Let us now give upper bounds on $|\mathcal{K}|$ and on $|U|$. Jordán [49, 50] showed that the size of the optimal augmenting edge set is at most $\max(b(G) - 1, \lceil \frac{t(G)}{2} \rceil) + \lceil \frac{k-2}{2} \rceil$. Here $b(G)$ is the maximum size of a clump, while $t(G)$ is the maximum number of pairwise disjoint sets in \mathcal{Q} .

Since $b(G) \leq n - (k - 1)$, $t(G) \leq n$, it follows that n is an upper bound on the size of an augmenting edge set. In a skeleton \mathcal{K} , the set of clumps connected by an edge xy forms a chain. Since the size of a chain can also be bounded by n , we may conclude $\sum_{X \in \mathcal{K}} (|K| - 1) \leq n^2$ and thus $|\mathcal{K}| \leq n^2$. Using the running time estimation in Section 3.5, this gives a bound $O(kn^5)$ on finding \mathcal{K} .

In Section 3.2.1 the minimum semi-cover of \mathcal{K} is reduced to a minimum symmetric chain cover of a poset $P = (U, \preceq, M)$ with $|U| = O(n^2)$, since there are $2|X|$ nodes in U corresponding to the clump $|X|$. Hence the running time of the matching algorithm may be bounded by $O(n^5)$. As indicated in the introduction, at most $\binom{n}{2}$ calls of the Dual Oracle enable us to compute an optimal augmentation. This gives a total running time $O(kn^7)$.

As in [36], another algorithm can be constructed which calls the dual oracle only once. First, let us find a skeleton $\mathcal{K} = \{K_1, \dots, K_\ell\}$ with a cover F and a grove Π of \mathcal{K} with $\text{def}(\Pi) = |F|$. Then we iteratively apply sequences of flipping operations as in Lemma 3.11 for $\mathcal{F}_{i-1} = \mathcal{C} \div \{K_1, \dots, K_{i-1}\}$ and K_i for $i = \ell, \ell - 1, \dots, 1$ resulting finally in a cover F' of \mathcal{C} with $|F| = |F'|$. For each i it can be easily seen that after $O(n^2)$ flippings we get a cover of \mathcal{F}_{i-1} , thus $O(n^4)$ improving flipping suffice. The realization of a flipping step can be done using similar techniques as in Section 3.5. We omit this analysis as it is highly technical and we could not get a better running time estimation as for the previous algorithm.

3.4 Further remarks

3.4.1 Node-induced cost functions

In this section, we show that the minimum cost version is also solvable for node-induced cost functions. $c' : E \rightarrow \mathbb{R}$ is a **node-induced cost function** if there exists a $c : V \rightarrow \mathbb{R}$ so that $c'(uv) = c(u) + c(v)$ for every $uv \in E$. By the second part of Theorem 3.12, for a skeleton \mathcal{K} and a node-induced cost function c' , the minimum c' -cost of a cover of \mathcal{C} is the same as that of \mathcal{K} . Hence it is enough to construct a subroutine for determining the minimum cost $\nu_{c'}(\mathcal{K})$ of a cover of \mathcal{K} . A minimum cost augmenting edge set can be found by iteratively calling this dual oracle.

Furthermore, by Lemma 3.14, $\nu_{c'}(\mathcal{K})$ equals the minimum cost of a semi-cover of \mathcal{K} . Finding a minimum-cost semi-cover can be easily done based on the following weighted version of Fleiner's theorem, which reduces to maximum cost matching in general graphs.

Given a symmetric poset $P = (U, \preceq, M)$ and a cost function $w : U \rightarrow \mathbb{R}$, let us define the cost of the symmetric chain $S = \{u_1v_1, \dots, u_\ell v_\ell\} \subseteq M$ with $u_1 \preceq \dots \preceq u_\ell$, $v_1 \succeq \dots \succeq v_\ell$ by $w(S) = w(u_\ell) + w(v_1)$. Our aim is now to find a chain cover of minimum total cost.

Consider the reduction to the matching problem in Section 3.3.2. For a matching $N \subseteq H$ of C , the components of $M \cup N$ are M -alternating paths each corresponding to a symmetric

chain. The alternating path corresponding to the chain S is $v_1u_1v_2u_2\dots v_\ell u_\ell$, hence the cost of the two nodes not covered by N equals the cost of the chain. Consequently, the cost of a symmetric chain cover equals the total cost of the nodes not covered by N . Hence minimizing the cost of a symmetric chain cover is equivalent to finding a maximum cost matching. Note that here we need a maximum cost matching only for node induced cost functions, although this can be found for arbitrary cost functions.

To find a minimum cost semi-cover of \mathcal{K} , we construct the symmetric poset $P = (U, \preceq', M)$ as in Section 3.2.1. For a one-way pair $u = (u^-, u^+) \in U$, let $w(u) = \min_{x \in u^+} c(x)$. We claim that finding a minimum cost symmetric chain cover for this w is equivalent to finding a minimum cost semi-cover of \mathcal{K} .

Indeed, there is a one-to-one correspondence between chains consisting of clumps of the form (X_i, X_i^*) and the symmetric chains of U (with the restriction that a chain may not contain both (X_i, X_i^*) , (X_j, X_j^*) for $i, j > 1$). A chain K_1, K_2, \dots, K_ℓ of clumps with orientations $L_1 \preceq L_2 \preceq \dots \preceq L_\ell$ can be covered by any edge between L_1^- and L_ℓ^+ , thus the minimum cost of an edge covering it is $w(L_\ell) + w(\overleftarrow{L_1})$ with w defined as above. Hence a minimum c -cost of a semi-cover in \mathcal{K} equals the minimum w -cost of a symmetric chain cover of P .

3.4.2 Degree sequences

What can we say about the degree sequences of the augmenting edge sets? It is well-known that in a graph G with some cost function on the edges, the sets of nodes covered by a minimum cost matching form the bases of a matroid. A natural generalization of matroid bases are base polytopes (see e.g. [69, Vol II, p. 767]).

For undirected edge-connectivity augmentation, the degree sequences of the augmenting edge sets form a base polytope, and the same holds for the in- and out-degree sequences for directed edge-connectivity augmentation (see e.g. [23]). This is also true in case of directed node-connectivity augmentation [31]. Moreover, all these results can be generalized for node-induced cost functions: the degree (resp. in- and out-degree) sequences of minimum cost augmenting edge sets form a base polytope. Hence a natural conjecture is the following:

Conjecture 3.27. *Given a $(k - 1)$ -connected graph G and a node-induced cost function, the degree sequences of minimum cost augmenting edge sets form a base polytope.*

This was essentially proved by Szabó proved in his master's thesis [70] for $k = n - 2$. His result holds even without the assumption that the graph is $(k - 1)$ -connected, indicating that the conjecture might hold for arbitrary graphs as well.

3.4.3 Abstract generalizations

In this section, we discuss possible generalizations and extension of the above results. A natural question is whether it is possible to give a generalization of Theorem 1.37 for abstract structures, in the sense as Theorem 1.2 generalizes Theorem 1.6 from strict one-way pairs in a $(k - 1)$ -connected graph to arbitrary crossing systems of set pairs. Indeed, it would be possible to formulate such an abstract theorem for describing coverings of a systems \mathcal{C} of “basic clumps”, where under basic clump we simply mean a subpartition of a set satisfying certain properties. However, it is not easy to extract the abstract properties \mathcal{C} needs to fulfill so that the argument carry over. In particular, we need to ensure Claim 3.1, Lemma 3.3, Claims 3.4 and 3.5, Lemma 3.20 and Lemma 2.5 (for set pairs arising from orientations of clumps). It may be verified that whenever \mathcal{C} satisfies these, all other proofs carry over; for the algorithm we also need a good representation of \mathcal{C} .

Since the argument is already quite abstract and complicated, and we could not find a short and nice list of properties that ensure all these claims, we did not formulate such an abstract theorem in order to avoid the addition of a new level of complexity. Furthermore, we believe that there should be a relatively simple abstract generalization of Theorem 1.37, which does not rely on all claims listed above. For comparison, the argument given in Chapter 2 for proving Theorem 1.6 strongly relies on properties of \mathcal{F} which hold only if \mathcal{F} is a crossing family of strict one-way pairs of a $(k - 1)$ -connected digraph (e.g. Claim 2.3, Lemma 2.5). Nevertheless, the more general Theorem 1.2 is true for arbitrary crossing families of set pairs, and admits a much simpler proof. (Recall that in Section 2.3.2 we also gave an extension of the “skeleton-proof” of Theorem 1.6 to that of Theorem 1.2 by introducing slim one-way pairs. Such an extension of Theorem 1.37 might also be possible, however, we would prefer a simpler type of argument.)

A natural application of such an abstract theorem would be rooted connectivity augmentation. Given a graph or digraph with designated node $r_0 \in V$, it is called **rooted k -connected** if there are at least k internally disjoint (directed) paths between r_0 and any other node. Similarly, a digraph is **rooted k -edge-connected** with root r_0 if there are at most $k - 1$ edge-disjoint directed path from r_0 to any other node. One may ask the augmentation questions for rooted connectivity as well. It turns out that for digraphs, the minimum cost versions of rooted k -connectivity and rooted k -edge-connectivity augmentation are both solvable in polynomial time (see Frank and Tardos [35] and Frank [29]): both problems can be formulated via matroid intersection (although the reduction of the node-connectivity version is far from trivial).

In contrast, for undirected graphs the minimum cost version of rooted k -connectivity augmentation is NP-complete: Hamiltonian cycle reduces to it even for $k = 2$ and 0-1 costs. The minimum cardinality version of augmenting rooted connectivity by one was studied by Nutov [68], who gave a an algorithm finding an augmenting edge set of size at most $opt + \min(opt, k)/2$.

An important difference between minimum cardinality directed and undirected rooted con-

nectivity augmentation is that while in the directed case there is an optimal augmenting edge set consisting only of edges outgoing from r_0 , in the undirected case it may contain edges not incident to r_0 . An example is $V = \{r_0, x, y, a\}$, $E = \{r_0x, r_0y, xa, ya\}$ (a rectangle). For $k = 3$, $F = \{xy, r_0a\}$ is an optimal augmenting set, but there is no augmenting set of size two of edges incident to r_0 .

We believe that a min-max formula and a polynomial time algorithm for finding an optimal solution could be given by extending the method of this chapter. However, it is not completely straightforward how clumps should be defined in this setting. At this point, we leave this question open, since we believe that it will be an easy consequence of a later general abstract theorem.

3.4.4 General connectivity augmentation

In what follows, we give an argument showing that there is no straightforward way of generalizing Theorem 1.37 for general connectivity augmentation. By "straightforward", we would mean a relation analogous to the one between Theorems 1.2 and 1.1: in the first one, the dual optimum is the maximum number of pairwise independent members of a crossing system of set pairs, while in the latter one, we are interested the maximum p -sum over pairwise independent set pairs. Hence a possible approach for general undirected connectivity augmentation would be the following. Let a clump be a subpartition $X = (X_1, \dots, X_\ell)$ of V with $d(X_i, X_j) = 0$ (we do not assume $|N_X| = k - 1$), and let $p(X)$ be a lower bound on the number of edges needed to cover X . There are multiple possible candidates for $p(X)$ and we do not commit to any of them, but work only with the natural assumption that $(\star) p(X) = \max(0, k - |N_X|)$ whenever $|X| = 2$; and $p(X) = 0$ whenever $|N_X| \geq k$. A natural conjecture is the following: the minimum size of an augmenting edge set equals the maximum deficiency of a grove, where in the definition of deficiency, each term $|X| - 1$ is replaced by $p(X)$.

We show by an example that this conjecture fails even if (\star) is the only assumption on $p(X)$. Let $G = (V, E)$ be the complement of the graph on Figure 3.3 and let $k = 9$. For a node $z \in V$, let $Z_z = (\{z\}, \{z\}^*)$. The only basic clumps in G with $|N_X| < 9$ are $Z_a, Z_b, Z_{u_1}, Z_{u_2}, Z_{v_1}, Z_{v_2}, (\{u_1, u_2\}, \{u_3\}, \{u_4\}), (\{v_1, v_2\}, \{v_3\}, \{v_4\})$ and $(\{a, c\}, \{b, d\})$. $\{u_1u_4, u_2u_3, v_1v_4, v_2v_3, ab, ad, bc\}$ is an augmenting edge set of size 7, while a grove of value 6 is the one consisting of two bushes $\mathcal{B}_1 = \{Z_{u_1}, Z_{u_2}, Z_{u_3}, Z_{u_4}, (\{a\}, \{u_1, u_2, d\})\}$ and $\mathcal{B}_2 = \{Z_{v_1}, Z_{v_2}, Z_{v_3}, Z_{v_4}, (\{b\}, \{v_1, v_2, c\})\}$.

We show that neither an augmenting edge set of size 6, nor a grove of value 7 exists. On the one hand, assume an augmenting edge set F exists with $|F| = 6$. Then F can be partitioned into $F = F_1 \cup F_2$ with $|F_1| = |F_2| = 3$, F_1 covering \mathcal{B}_1 and F_2 covering \mathcal{B}_2 . However, we need at least two edges to cover Z_a and two to cover Z_b , and these can only be contained in F_1 and F_2 , respectively. If $ad \in F_1$, then F_1 cannot contain any of au_1 and au_2 as otherwise at least one of Z_{u_3} and Z_{u_4} would remain uncovered. Hence $ad \notin F_1$, and similarly $bc \notin F_2$. $ab, cd \notin F$ as they




Figure 3.3: Example concerning general connectivity augmentation.

do not cover any of \mathcal{B}_1 and \mathcal{B}_2 , thus $(\{a, c\}, \{b, d\})$ remains uncovered.

On the other hand, assume a grove of value 7 exists. We claim that it should contain $(\{a, c\}, \{b, d\})$, and two clumps of the form $(\{a\}, A)$ and $(\{b\}, B)$ with $b \in A$ and $a \in B$. This is clearly a contradiction as they cannot be simultaneously contained in a grove, since the edge ab connects all three of them. It can easily be checked that if we do not require $(\{a, c\}, \{b, d\})$ to be covered, then the remaining clumps may all be covered by six edges. The same holds unless we require all clumps of the form $(\{a\}, A)$ with $b \in A$ and all clumps of $(\{b\}, B)$ with $a \in B$ to be covered. Consequently, every grove of value 7 should contain such clumps.

3.5 Implementation via bipartite matching

In this section we present how the subroutine for constructing a skeleton can be implemented using bipartite matching theory. The argument follows the same lines as the one in Section 2.4; we adopt the terminology, notation and multiple fundamental claims proved there. Before starting the reduction to bipartite graphs, let us prove a simple claim concerning pieces. This is an analogue of Claim 2.3.

Claim 3.28. *For a piece $Y \in \mathcal{Q}_1$ and an arbitrary set $X \subseteq V$, if $X^* \supseteq Y^*$, then $X \subseteq Y$.*

Proof. Indeed, assume X is not a subset of Y , thus $|X \cup Y| > |Y|$. The condition gives $(X \cup Y)^* = Y^*$, and hence $|N(X \cup Y)| < |N(Y)| = k - 1$, contradicting that G is $(k - 1)$ -connected. \square

Given the $(k - 1)$ -connected graph $G = (V, E)$, let us construct the bipartite graph $B = (V', V''; H)$ as follows. With each node $v \in V$ associate nodes $v' \in V'$ and $v'' \in V''$ and an

edge $v'v'' \in H$. With each edge $uv \in E$ associate two edges $v'u'', u'v'' \in H$. For a set $X \subseteq V$, we denote by X' and X'' its images in V' and V'' , respectively. The $(k-1)$ -connectivity of G implies that B is a $(k-1)$ -elementary bipartite graph. For a set $X \subseteq V$, X' is tight if and only if $X \in \mathcal{Q}_1$. (Recall that in Section 2.4 we called a set $X' \subseteq V'$ tight if $|\Gamma(X')| = |X'| + k - 1$ and $\Gamma(X') \neq V''$.)

First we need to find the set \mathcal{M} of minimal pieces. This is done by computing the edge set F_{uv} (the $(k-1)$ - uv -factor) by a single max-flow computation for every $u, v \in V$, $uv \notin E$. By Claim 2.15, the minimal $u\hat{v}$ -sets can be found by a breadth-first search. The minimal ones among these will give the elements of \mathcal{M} .

Consider now a stable cross-free \mathcal{H} which is not complete, a minimal element $M \in \mathcal{M} - \bigcup \mathcal{H}$ and $\mathcal{L}_1, \mathcal{L}_2$ as defined by (3.2). If $\mathcal{L}_1 = \emptyset$ then we are done by Claim 3.23, hence in the sequel we assume $\mathcal{L}_1 \neq \emptyset$.

By Lemma 3.24, our task is to find a minimal C fitting (\mathcal{H}, M) . Let \mathcal{T} be the set of the maximal ones among those pieces of the clumps in \mathcal{L}_1 which are subsets of M^* .

Claim 3.29. *\mathcal{T} consists of pairwise disjoint sets.*

Proof. Consider clumps $X, Y \in \mathcal{L}_1$ with pieces $X_1, Y_1 \in \mathcal{T}$. If X and Y are independent then $X_1 \cap Y_1 = \emptyset$ as otherwise an edge between $X_1 \cap Y_1$ and M would connect both. If they are dependent, then we show that the dominant side X_i of X w.r.t Y is different from X_1 . Indeed, if $X_i = X_1$, then the dominant side of Y w.r.t. X should be $Y_j \neq Y_1$ as otherwise $M \subseteq Y_1$ would follow. Hence $Y_1 \subsetneq X_1$, a contradiction to the maximality of Y_1 . Similarly, the dominant side of Y w.r.t. X may not be Y_1 . Hence $Y_1 \subseteq X^*$, thus $X_1 \cap Y_1 = \emptyset$. \square

Let us construct the bipartite graph $B_1 = (V', V''; H_1)$ from B by adding some new edges as follows. (1) For each $X \in \mathcal{L}_2$, let $x'y'', y'x'' \in H_1$ for every xy connecting X . (2) Let $x'y'' \in H_1$ whenever $T \in \mathcal{T}$, $x \in T$ and $y \in T \cup N(T)$. (3) For each $X \in \mathcal{L}_1$ with dominant piece X_a w.r.t. M^\sharp , let $x'y'' \in H_1$ for every $x \in X_a, y \in X_a^*$.

Claim 3.30. *Let $C \in \mathcal{Q} - \bigcup \mathcal{H}$, $C \subseteq M^*$, supported by some $W \in \mathcal{H}$. C fits (\mathcal{H}, M) if and only if C' is tight in B_1 .*

Proof. $C' \subseteq V'$ is tight in B_1 if and only if it is tight in B and there is no edge in $x'y'' \in H_1 - H$ with $x' \in C', y' \in V'' - \Gamma(C')$ (or equivalently, xy connects the clump (C, C^*)).

Assume C fits (\mathcal{H}, M) . Property (d) forbids that any $x'y'' \in H_1 - H$ of the first type cover C' , while (c) forbids any $x'y''$ of the second or third type to cover C' . For the other direction, properties (a) and (b) follow by the conditions. For (d), if C were dependent with some $X \in \mathcal{L}_2$, then a new edge of the first type would cover C' . For (c), if $C \cap X_i \neq \emptyset, X_i - C \neq \emptyset$ for some $X \in \mathcal{L}_1$ with a piece $X_i \subsetneq M^*$, then consider a $T \in \mathcal{T}$ with $X_i \subseteq T$. $C - T \neq \emptyset$ as otherwise W, C^\sharp, T^\sharp would contradict stability. By Claim 3.28, $C^* \cap (T \cup N(T)) \neq \emptyset$, hence a new edge of

the second type covers C' . Finally, if X_a is the dominant piece of X w.r.t. $M^\#$ and $X_a \cap C \neq \emptyset$, $X_i \cap C^* \neq \emptyset$, then there is a new edge of the third type covering C' . \square

To find a C as in Lemma 3.24, we need to add some further edges to B_1 . Indeed, we need to ensure that $C \in \mathcal{Q} - \bigcup \mathcal{H}$ and furthermore that C is supported by some $W \in \mathcal{L}_1$. Consider now a $W \in \mathcal{L}_1$ with a piece $W_1 \in \mathcal{T}$ and a connected set Q with $W_1 \subsetneq Q \subseteq M^*$. Let $Z(Q)$ denote the unique minimal X satisfying the following property:

$$X \in \mathcal{Q}, Q \subseteq X, \text{ and } X \text{ fits } (\mathcal{H}, M). \quad (3.3)$$

We will determine $Z(Q)$ for different sets Q in order to find K . As in the directed case, it is easy to see that $Z(Q)$ is well-defined. The next claim gives an easy algorithm for finding $Z(Q)$ for a given Q .

Claim 3.31. *Fix some $u \in Q$, $v \in M$. Let B_2 denote the graph obtained from B_2 by adding all edges $u'y''$ with $y \in Q \cup N(Q)$. Let S denote the set of nodes z for which there exists an alternating path for F_{uv} from u' to z' . Then $Z(Q) = S$.*

Proof. As M^* is an $u\hat{v}$ -set in B_2 , applying Claim 2.15(a) for B_2 instead of B , we get that B_2 contains no alternating path for F_{uv} between u' and v'' . By Claim 2.15(b), S is the unique minimal uv -piece in B_2 . $\Gamma(S' \cup Q') = \Gamma(S')$ thus $Q \cup N(Q) = S \cup N(S)$ because of the new edges in B_2 , hence by Claim 3.28, $Q \subseteq S$. By making use of Claim 3.30, S is the unique minimal set satisfying (3.3), thus $Z(Q) = S$. \square

Consider now a clump $W = (W_1, W_2, \dots, W_h) \in \mathcal{L}_1$ with $W_1 \in \mathcal{T}$. We want to find a Z_W fitting (\mathcal{H}, M) supported by W_1 . For each $q \in N_W \cap M^*$, let us compute $Z(Q)$ for $Q = W + q$. Let C_W denote a minimal set among these. A $Z(Q)$ can be found by a single breadth-first search, thus we need at most $k - 1$ breadth-first searches. We may compute such a C_W for all possible choices of W , and a minimal among these gives a minimal C fitting (\mathcal{H}, M) . Therefore the running time may be bounded by $(k - 1)n$ breadth-first searches since by Claim 3.29, $|\mathcal{T}| \leq n$. Somewhat surprisingly, this better compared to the directed case, where we needed n^2 breadth first searches. The reason is that here we could take advantage of the fact that all pieces in a basic clump are connected and therefore consider only $Q = W + q$ for $q \in N_W \cap M^*$. In contrast, the tail or a head of a one-way pair may contain a directed cut and therefore we had to examine a larger set of Q 's.

Complexity

To find a skeleton system first we need n^2 Max Flow computations to determine the minimal pieces and the auxiliary graphs. The running time for extending the stable cross-free system by one member is dominated by $(k - 1)n$ breadth first searches. Thus if s is an upper bound

on the size of a skeleton, then we can determine one in $O(n^5 + skn^3)$ running time by using an $O(n^3)$ maximum flow algorithm and an $O(n^2)$ breadth first search algorithm.

Chapter 4

General directed node-connectivity augmentation

The results in this chapter were published in [74], a joint paper with András Benczúr jr. We have defined posets with the strong interval property and formulated Theorem 1.40 in Section 1.5.2. Let us start with the proof of Claim 1.39.

Proof of Claim 1.39. Property (i) of Definition 1.38 follows directly by the properties of set union, intersection and containment. The relation between intervals and subfamilies defined by pairs of nodes is straightforward since the minimal elements of \mathcal{S} are the set pairs of the form $(\{u\}, V - u)$ and the maximal ones are of the form $(V - v, \{v\})$. To prove Property (ii), consider an edge xy with $[m, M] = I_{xy}$. (1.7) is a consequence of Lemma 2.2. \square

We have already seen that Theorem 1.1 follows from Theorem 1.40. Let us now show that the reverse implication also holds and hence they are equivalent. Given a poset (\mathcal{P}, \preceq) with the strong interval property, let us define a representative element $\varphi(x)$ for every minimal or maximal element x . For $a \in \mathcal{P}$, let us define the pair $\Psi(a) = (a^-, a^+)$ so that

$$a^- = \{\varphi(m) : m \preceq a, m \in \mathcal{P} \text{ minimal}\}; \quad a^+ = \{\varphi(M) : M \succeq a, M \in \mathcal{P} \text{ maximal}\}.$$

It is easy to show that the function Ψ is a homomorphism for \vee , \wedge and \preceq . Let us define $p'(K) := \max\{p(a) : \Psi(a) = K\}$ where $p'(K) = 0$ if there exists no $a \in \mathcal{P}$ with $\Psi(a) = K$. It is easy to verify that this is positively crossing supermodular. Hence applying Theorem 1.1 for p' on the set pairs implies Theorem 1.40.

Let us now show some basic properties of the tight elements.

Lemma 4.1. *If x and y are two dependent tight elements with $p(x) > 0$, $p(y) > 0$, then both $x \vee y$ and $x \wedge y$ are tight.*

Proof. Let $g(x)$ denote the number of intervals covering element x . By the strong interval property all intervals that cover $x \vee y$ or $x \wedge y$ also cover x or y and if they cover both, then they cover all four, hence $g(x) + g(y) \geq g(x \vee y) + g(x \wedge y)$. The proof is complete by

$$\begin{aligned} g(x \vee y) + g(x \wedge y) &\geq p(x \vee y) + p(x \wedge y) \geq \\ &\geq p(x) + p(y) = g(x) + g(y) \geq g(x \vee y) + g(x \wedge y) \end{aligned} \tag{4.1}$$

implying equality everywhere. Here the first inequality follows since we have a cover; the second is the definition of crossing supermodularity; and the equality follows by the tightness of x and y . \square

The following easy corollary will be used throughout the paper:

Corollary 4.2. *For a cover \mathcal{I} , every $I \in \mathcal{I}$ has a unique minimal and a unique maximal tight element.*

Lemma 4.3. *If x and y are two dependent tight elements with $p(x) > 0$, $p(y) > 0$, and the interval $[m, M] \in \mathcal{I}$ contains x , then it contains at least one of $x \vee y$ and $x \wedge y$; or equivalently, $y \preceq M$ or $m \preceq y$.*

Proof. Recall that by the proof of Lemma 4.1 we have equality everywhere in (4.1); the last inequality hence turns to $g(x) + g(y) = g(x \vee y) + g(x \wedge y)$. By the strong interval property all intervals that cover $x \vee y$ or $x \wedge y$ also cover x or y and if they cover both, then they cover all four. Hence the above equality implies the claim. \square

4.1 The algorithm

We give a brief overview of our algorithm for the 0–1 valued case (Theorem 1.2) first. The algorithm starts out with a (possible greedy) interval cover $\mathcal{I} = \{I_1, \dots, I_r\}$. In Algorithm PUSH-DOWN-REDUCE we maintain a tight element $u_i \in I_i$ for each interval I_i as a witness for the necessity of I_i in the cover. As long as the set of witnesses are non-independent, in Procedure PUSHDOWN we replace certain u_i by smaller elements. By such steps we aim to arrive in an independent system of witnesses. If witnesses are indeed pairwise independent, they form a dual solution with the same value as the primal cover solution, thus showing both primal and dual optimality. Otherwise in Procedure PUSHDOWN the Procedure REDUCE is called, a procedure that exchanges interval endpoints so that we get an interval cover of size one less.

In order to handle weighted posets, technically we need to consider multisets of intervals and witnesses in our algorithm. We assume $\mathcal{I} = \{I_1, \dots, I_r\}$ may contain the same interval more than once and the same may happen to the set of witnesses. The next lemma shows that if the witnesses are pairwise independent *as a weighted set* instead of a multiset, then the solution is optimal.

Lemma 4.4. Consider a cover $\mathcal{I} = \{I_1, \dots, I_r\}$ and a tight element $u_i \in I_i$ for every i . If for every i, j , u_i and u_j are either independent or $u_i = u_j$, then the elements $\{u_1, \dots, u_r\}$ give a dual optimal solution, and hence \mathcal{I} is an optimal cover.

Proof. It suffices to show that if for a poset element y there exists an i with $y = u_i$, then there exist exactly $p(y)$ such intervals I_j with $y = u_j$. Since $y = u_i$ is tight, there are exactly $p(y)$ intervals I_j with $y \in I_j$. Consider such an u_j now: u_i and u_j are either independent or $u_i = u_j$, but the first case is impossible since both of them are covered by I_j . Hence $u_j = u_i$ for all $p(y)$ values of j . \square

Algorithm PUSHDOWN-REDUCE(\mathcal{I})

for $j = 1, \dots, r$ **do**

if I_j has no tight elements **then**

return reduced cover $\{I_i : i = 1, \dots, j - 1, j + 1, \dots, r\}$

$u_j^{(1)} \leftarrow$ maximal tight element of I_j

$t \leftarrow 1$

do

for $j = 1, \dots, r$ **do**

$u_j^{(t+1)} \leftarrow$ PUSHDOWN(j, t, \mathcal{I})

$t \leftarrow t + 1$

while exist j such that $u_j^{(t)} < u_j^{(t-1)}$

return dual optimal solution $\{u_1^{(t)}, \dots, u_r^{(t)}\}$

Procedure PUSHDOWN(j, t, \mathcal{I})

$U \leftarrow \{x : m_j \preceq x \preceq u_j^{(t)}, x \text{ tight and } \forall i = 1, \dots, r, u_i^{(t)} \text{ may not push } x \text{ down}\}$

if $U = \emptyset$ **then**

$t^* \leftarrow t;$

return REDUCE(j, t^*, \mathcal{I})

else return the maximal $x \in U$

4.1.1 The PUSHDOWN step

Our Algorithm PUSHDOWN-REDUCE (see box) tries to push witnesses down along their intervals in iterations $t = 1, 2, \dots$ until they satisfy the requirements of Lemma 4.4; witnesses are




Figure 4.1: Different cases when u may push v down. By Lemma 4.3 $m_i \preceq v$, and there are three possible cases: (a) $m_j \not\preceq u \preceq M_j$, (b) $m_j \preceq u \not\preceq M_j$, and (c) $m_j \preceq u \preceq M_j$

superscripted by the iteration value (t). Initial witnesses $u_j^{(1)}$ are maximum tight; their existence follows by Corollary 4.2.

Given two intervals $I_i = [m_i, M_i]$ and $I_j = [m_j, M_j]$ and two tight elements $u \in I_i$ and $v \in I_j$, we say that u may **push v down** with respect to I_i if u and v are dependent and $v \not\preceq M_i$. In the case set U of Procedure **PUSHDOWN** (see box) is nonempty we will push v down, i.e. replace it by the maximal element of U strictly below v . Notice that the definition depends on the choice of the interval I_i with $u \in I_i$; it is possible that v may push u down with respect to certain I_i and not with others. In the following, when it is clear from the context, we will omit mentioning I_i . Different scenarios when u may push v down are shown in Figure 4.1.

In what follows we motivate which element replaces a given v when v gets pushed down. When selecting $u_j^{(t+1)}$, our aim is to replace $u_j^{(t)}$ by the maximal such tight element $x \in I_j$ which satisfies $x \preceq u_j^{(t)}$ and no $u_i^{(t)}$ may push x down. As the motivation of pushing $u_j^{(t)}$ down by $u_i^{(t)}$ we give the following claim as a relatively easy consequence of Lemma 4.9; we omit the proof as it is not used elsewhere. If $u_i^{(t)}$ may push $u_j^{(t)}$ down, then for all subsequent $t' > t$ of the **WHILE** loop of Algorithm **PUSHDOWN-REDUCE** if the witnesses $u_j^{(t')}$ and $u_i^{(t')}$ are dependent then they must be equal. This will be the main reason why all non-equal dependent pairs of witnesses gradually disappear from the system.

While the above motivation considers the dual solution, namely it shows that the set of witnesses will satisfy the optimality requirements, we may also give a primal motivation of pushing v down by u . If u is maximum tight in I_i , then we may hope that by replacing $[m_i, M_i]$ by $[m_i, M_j]$ we still get a cover. In the examples of Figure 4.1 this holds for cases (a) and (c). In this cover v is contained in the new interval while it was not contained in the old, thus it may be replaced by a smaller witness.

However, this argument fails for case (b) since $u \notin [m_i, M_j]$ and the actual proof of correctness will use a slightly more complicated argument. In the case of increasing connectivity by

one (see [4]), the only possible scenario was (a). This is the main reason why the analysis is significantly harder for the general case. While the argument for replacing $[m_i, M_i]$ by $[m_i, M_j]$ fails, we still push v down and proceed with the algorithm. Then we use a backward analysis as in [4]; in the weighted case it turns out that, while this fails to hold in general, if a particular interval exchange is performed corresponding to a pushdown step, then the exchange is valid and in particular we have $u \preceq M_j$. We prove this later in Lemma 4.12.

The next properties of elements that one may push the other down are required both for the definition of the algorithm and later for the proof of correctness.

Lemma 4.5. *If $u, u' \in I_i$ and $v \in I_j$ are tight with $u' \preceq u$ and u may push v down, then u' may also push v down.*

Proof. We only have to show that u' and v are dependent. $v \not\preceq M_i$, since u may push v down. Now by Lemma 4.3 we have $m_i \preceq v$. Hence the dependence of u' and v follows: a common lower bound is m_i and a common upper bound is $u \vee v$. \square

Lemma 4.6. *Suppose $u \in I_i$, $v \in I_j$, $v' \in I_h$ are tight elements and v and v' are dependent. If u may push $v \vee v'$ down, then it may also push either v or v' down.*

Proof. Since u may push $v \vee v'$ down, we have $v \vee v' \not\preceq M_i$, hence $m_i \preceq v \vee v'$ by Lemma 4.3. By the strong interval property either $m_i \preceq v$ or $m_i \preceq v'$. By symmetry let us consider the first case; in this case v and u are also dependent since their common lower bound is m_i and their common upper bound is $u \vee (v \vee v')$. If $v \not\preceq M_i$, then u may push v down. Suppose now $m_i \preceq v \preceq M_i$. Since u may push $v \vee v'$ down, we have $v \vee v' \not\preceq M_i$ and thus $v' \not\preceq M_i$. Then by applying Lemma 4.3 for v, v' and $[m_i, M_i]$ it follows that $m_i \preceq v'$, hence u and v' are dependent. Finally by $v' \not\preceq M_i$ we get that u may push v' down. \square

The actual change of a witness $u_j^{(t)}$ is performed in Procedure PUSHDOWN (see box). We select all tight elements $x \in I_j$, $x \preceq u_j^{(t)}$ into a set U that cannot be pushed down with elements $u_i^{(t)}$. If U is nonempty, we next show that it has a unique maximal element; we use this element as the new witness $u_j^{(t+1)}$.

Lemma 4.7. *In Procedure PUSHDOWN either $U = \emptyset$ or else it has a unique maximal element.*

Proof. It suffices to show that if $x, x' \in V$, then so is $x \vee x' \in V$. Obviously, $x \vee x'$ is tight and $m_j \preceq x \vee x' \preceq u_j^{(t)}$. Suppose now that some $u_i^{(t)}$ may push $x \vee x'$ down. By Lemma 4.6, $u_i^{(t)}$ may push either x or x' down, contradicting $x, x' \in U$. \square

If we find no dependent pair of witnesses such that one may push the other down, then we will show that the witnesses are pairwise independent or equal and thus the solution is optimal. As long as we find pairs such that one may push the other down, in the main loop of Algorithm PUSHDOWN-REDUCE we record a possible interval endpoint change by pushing one witness lower in its interval; these changes are then unwound to a smaller cover as shown in Section 4.1.3.

4.1.2 Proof for termination without REDUCE

We turn to the first key step in proving the correctness: we show that if the algorithm terminates without calling Procedure REDUCE, then $u_i^{(t)}$ are pairwise independent or equal; in other words, if none of them may be pushed down by another, then the solution is optimal.

Theorem 4.8. *If the algorithm terminates without calling Procedure REDUCE, then $u_i^{(t)}$ and $u_j^{(t)}$ dependent implies $u_i^{(t)} = u_j^{(t)}$.*

The theorem is an immediate consequence of the next lemma. To see, notice that if the algorithm terminates without calling Procedure REDUCE, then in a last iteration the **while** condition of Algorithm PUSHDOWN-REDUCE fails. However then there are no pairs i and j such that $u_i^{(t)}$ may push $u_j^{(t)}$ down.

Lemma 4.9. *Assume that $t_1 \leq t_2$, and $u_i^{(t_2)}$ and $u_j^{(t_1)}$ are dependent, and $u_j^{(t_1)}$ may not push $u_i^{(t_2)}$ down. Then $u_i^{(t_2)} \preceq u_j^{(t_1)}$.*

This lemma is used not only for proving Theorem 4.8 but also in showing the correctness of Procedure REDUCE in Section 4.1.3 via the next immediate corollary.

Corollary 4.10. *If $u_j^{(t)}$ and $u_i^{(t+1)}$ are dependent, then $u_i^{(t+1)} \preceq u_j^{(t)}$.*

In the proof of Lemma 4.9 we need to characterize elements that cause witness u_j move below a certain tight element y . Assume that for some tight $y \in I_j$ and t we have $y \not\preceq u_j^{(t)}$. Since $u_j^{(1)}$ is maximal tight, we may select the unique t_0 with $y \preceq u_j^{(t_0)}$ but $y \not\preceq u_j^{(t_0+1)}$. In step PUSHDOWN(j, t_0, \mathcal{I}) we must have an $u_d^{(t_0)}$ that may push y down. We will use this in the following special case:

Lemma 4.11. *Assume that z is tight and dependent from $u_j^{(t)}$. Assume furthermore that $z \not\preceq u_j^{(t)}$ and $z \preceq M_j$. Then there exists $t_0 < t$ and d such that $u_d^{(t_0)}$ may push $u_j^{(t)} \vee z$ down. In addition, $u_d^{(t_0)}$ may also push z down.*

Proof. We apply the above observations for $y = u_j^{(t)} \vee z \in I_j$. Since y is tight, $y \preceq u_j^{(1)}$. And since $z \not\preceq u_j^{(t)}$, we get $y = u_j^{(t)} \vee z \not\preceq u_j^{(t)}$. We select t_0 with $y \preceq u_j^{(t_0)}$ but $y \not\preceq u_j^{(t_0+1)}$; then in step PUSHDOWN(j, t_0, \mathcal{I}) we must have an $u_d^{(t_0)}$ that may push y down.

For the second part of the claim observe that by Lemma 4.6, $u_d^{(t_0)}$ may push either $u_j^{(t)}$ or z down. The first choice is impossible, since then $u_d^{(t-1)}$ could also push $u_j^{(t)}$ down by Lemma 4.5, and $t - 1 \geq t_0$. This latter contradicts the choice of $u_j^{(t)}$ as the maximum tight element that may not be pushed down in PUSHDOWN($j, t - 1, \mathcal{I}$). \square

Proof of Lemma 4.9. $u_i^{(t_2)} \preceq M_j$, since $u_j^{(t_1)}$ may not push $u_i^{(t_2)}$ down. If $u_i^{(t_2)} \not\preceq u_j^{(t_1)}$, then the conditions of Lemma 4.11 hold with $z = u_i^{(t_2)}$ and $t = t_1$. Thus we have some $t_0 < t_1$ and d such




Figure 4.2: Procedure REDUCE called with $t^* = 1$. The two upright intervals are the original ones with their tight elements shaded. These two intervals will be replaced by the single bold interval. The new interval contains all tight elements of the old ones since $u_{j_2}^{(1)} \preceq M_{j_1}$ by Lemma 4.12. Remember that the intervals need not to be disjoint.

that $u_d^{(t_0)}$ may push $z = u_i^{(t_2)}$ down. But then $u_d^{(t_2-1)}$ may also push $u_i^{(t_2)}$ down by Lemma 4.5. This latter contradicts the choice of $u_i^{(t_2)}$ as the maximum tight element that may not be pushed down in $\text{PUSHDOWN}(i, t_2 - 1, \mathcal{I})$. \square

4.1.3 The REDUCE step

So far we have proved that if REDUCE is not called, then the initial primal solution is optimal and the algorithm finds a dual optimum proof of this fact. Now we turn to the second scenario when Procedure REDUCE is called; in this case the solution is not optimal, since Procedure REDUCE is called from Procedure PUSHDOWN when $U = \emptyset$. This means $u_j^{(t)} \notin U$ and thus there exists an i such that $u_i^{(t)}$ may push $u_j^{(t)}$ down.

Procedure REDUCE is called when one witness disappears from the dual solution. In this case we unwind the steps to find a cover of size one less in Procedure REDUCE based on interval exchanges at certain pairs of tight poset elements.

To illustrate the idea of Procedure REDUCE, first we discuss the simplest case $t^* = 1$; the general case will then be reduced to this case by a special induction. We summarize Procedure REDUCE-ONESTEP for this particular scenario with steps shown in Figure 4.2. Since $t^* = 1$, we have some $1 \leq j_1 \leq k$ such that Procedure REDUCE is called within Procedure PUSH-

Procedure REDUCE-ONESTEP(j, \mathcal{I})

$j_1 \leftarrow j$;

$q \leftarrow$ minimal tight element in $[m_{j_1}, M_{j_1}]$

$j_2 \leftarrow$ **minimum** value $\ell \neq j_1$ such that $u_\ell^{(1)}$ may push q down

return reduced cover $\{[m_i, M_i] : 1 \leq i \leq r, i \neq j_1, j_2\} \cup \{[m_{j_2}, M_{j_1}]\}$.

DOWN($j_1, 1, \mathcal{I}$). This means that

$$U = \{x : m_{j_1} \preceq x \preceq u_{j_1}^{(1)}, x \text{ tight and } \forall \ell = 1, \dots, r, u_\ell^{(1)} \text{ may not push } x \text{ down}\}$$

is empty. By Corollary 4.2, $[m_{j_1}, M_{j_1}]$ has a unique minimal tight element q ; since $q \notin U$, we must have some $\ell = j_2$ such that $u_\ell^{(1)}$ may push q down. Given an ordering over the intervals, the algorithm selects j_2 as the minimal such ℓ and returns a reduced interval system

$$\mathcal{I} - [m_{j_1}, M_{j_1}] - [m_{j_2}, M_{j_2}] + [m_{j_2}, M_{j_1}]. \quad (4.2)$$

In the proof of case $t^* = 1$ we use the following general lemma for $h = j_1$, $\ell = j_2$, $u = u_{j_2}^{(1)}$.

Lemma 4.12. *Let q be the minimal tight element of I_h . If $u \in I_\ell$ may push q down, then $u \preceq M_h$. Furthermore for all tight $v \in I_h$ we have that u may push v down with respect to I_ℓ .*

Proof. Suppose by contradiction that $u \not\preceq M_h$. Since u and q are dependent, by Lemma 4.3, $u \wedge q \in I_h$. Since q is the minimal tight in I_h , we have $q \preceq u \wedge q$, hence $q \preceq u \preceq M_\ell$, contradicting that u may push q down. For the second part of the claim, consider a tight element $v \in I_h$. Elements u and v are dependent, since common lower and upper bounds are $u \wedge q$ and M_h , respectively. By $q \preceq v$ and $q \not\preceq M_\ell$ the required $v \not\preceq M_\ell$ follows. \square

Lemma 4.13. *If $t^* = 1$, Procedure REDUCE-ONESTEP(j_1, \mathcal{I}) returns an interval cover.*

Proof. It suffices to show that $[m_{j_2}, M_{j_1}]$ contains all tight elements of both $[m_{j_1}, M_{j_1}]$ and $[m_{j_2}, M_{j_2}]$; furthermore there is no common tight element in $[m_{j_1}, M_{j_1}]$ and $[m_{j_2}, M_{j_2}]$. In this case we may replace the intervals $[m_{j_1}, M_{j_1}]$ and $[m_{j_2}, M_{j_2}]$ by $[m_{j_2}, M_{j_1}]$ since if a tight element is contained by exactly one of $[m_{j_1}, M_{j_1}]$ and $[m_{j_2}, M_{j_2}]$ then it is contained by the new interval and containment by both is excluded.

To prove, first let $x \in [m_{j_2}, M_{j_2}]$ be tight; $x \leq u_{j_2}^{(1)}$ by maximality. When applying Lemma 4.12 for $h = j_1$, $\ell = j_2$, $u = u_{j_2}^{(1)}$, we get $u_{j_2}^{(1)} \preceq M_{j_1}$. This implies $m_{j_2} \preceq x \preceq u_{j_2}^{(1)} \preceq M_{j_1}$, as required.

Next let $x \in [m_{j_1}, M_{j_1}]$ be tight; $q \preceq x$ for the minimal tight q of $[m_{j_1}, M_{j_1}]$. By Lemma 4.3, $m_{j_2} \preceq q$, thus we get $m_{j_2} \preceq q \preceq x \preceq M_{j_1}$ as required.

Finally assume that a common tight element $x \in [m_{j_1}, M_{j_1}] \cap [m_{j_2}, M_{j_2}]$ exists; now $q \preceq x \preceq M_{j_2}$, contradicting the fact that $u_{j_2}^{(1)}$ may push q down. \square

Procedure REDUCE(j, t^*, \mathcal{I})

$j_1 \leftarrow j$;
for $t = t^*, \dots, 1$ **do**
 $s \leftarrow t^* + 1 - t$
 $q \leftarrow$ minimal tight element in $[m_{j_s}, M_{j_s}]$
 $j_{s+1} \leftarrow$ **minimum** value $\ell \neq j_s$ such that $u_\ell^{(t)}$ may push q down
 $m_{j_s} \leftarrow m_{j_{s+1}}$
return reduced cover $\{[m_i, M_i] : 1 \leq i \leq r, i \neq j_{t^*+1}\}$.

Our aim in Procedure REDUCE (see box) is to repeatedly pick an interval $[m_{j_s}, M_{j_s}]$ and try to find another interval $[m_{j_{s+1}}, M_{j_{s+1}}]$ such that if we replace $[m_{j_s}, M_{j_s}]$ by $[m_{j_{s+1}}, M_{j_s}]$, then after the switch the minimum tight element of $[m_{j_{s+1}}, M_{j_{s+1}}]$ increases. We ensure this by defining

$$j_{s+1} \leftarrow \text{minimum value } \ell \neq j_s \text{ such that } u_\ell^{(t)} \text{ may push } q \text{ down,}$$

where q is the minimum tight element of $[m_{j_s}, M_{j_s}]$ after the interval changes and $t = t^* + 1 - s$. Applying Lemma 4.12 for $h = j_s$, $\ell = j_{s+1}$, $u = u_{j_{s+1}}^{(t)}$ we get $u_{j_{s+1}}^{(t)} \leq M_{j_s}$. Thus when replacing $[m_{j_s}, M_{j_s}]$ by $[m_{j_{s+1}}, M_{j_s}]$, the tight elements x in $[m_{j_{s+1}}, M_{j_{s+1}}]$ with $x \leq u_{j_{s+1}}^{(t)}$ will no longer be tight after the switch. The overall idea is seen in Figure 4.3.

While the first step of the procedure is well-defined since we call Procedure REDUCE exactly when the minimal tight $q \in I_j$ for $j = j_1$ is pushed down by certain other $u_\ell^{(t^*)}$, the existence of such an ℓ is by no means obvious for all the other iterations of the main loop as switches among the intervals could completely rearrange the set of the tight elements.

The existence of all further ℓ in Procedure REDUCE as well as the correctness of the algorithm is proved by “rewinding” the algorithm after the first iteration of Procedure REDUCE and showing that each step is repeated identical up to iteration $t^* - 1$. The intuition behind rewinding is based on the resemblance of Procedure REDUCE to an augmenting path algorithm. In this terminology, instead of directly proving augmenting path properties we use a special induction by executing the main loop of the procedure step by step and after each iteration rewinding the main algorithm. In the analogy of network flow algorithms, this may correspond to analyzing an augmenting path algorithm by choosing path edges starting at the source, changing the flow along this edge to a preflow, and at each step proving that the remaining path augments the flow.

The key Theorem below will show, by induction on the value t^* of t at the termination of the main loop of Algorithm PUSHDOWN-REDUCE, that the intermediate modified interval sets are covers for $t^*, t^* - 1, \dots, 1$. Finally when applied for $t^* = 1$ we get that Procedure REDUCE finds an interval cover of size one less than before by Lemma 4.13. This completes the correctness analysis of Procedure REDUCE. Before stating the Theorem, we define the intermediate modified




Figure 4.3: Procedure REDUCE called with $t^* = 2$. The three upright intervals are the original ones with their tight elements shaded. The original three intervals will be replaced by the two bold intervals using the marked witnesses. Note that the two new intervals contain all tight elements of the old ones. While the number of intervals covering certain non-tight elements (x in the example) may decrease, we prove that they remain covered. Note that the original intervals are not necessarily disjoint.

interval set \mathcal{I}' and show it is a cover.

Lemma 4.14. *Let*

$$\mathcal{I}' = \mathcal{I} - [m_{j_1}, M_{j_1}] + [m_{j_2}, M_{j_1}]. \quad (4.3)$$

be the set of intervals after the first iteration of Procedure REDUCE. Then \mathcal{I}' is a cover.

Proof. Since $u_{j_2}^{(1)}$ may push q down, $q \not\leq M_{j_2}$, thus by Claim 4.3, $m_{j_2} \preceq q$ and so $[m_{j_2}, M_{j_1}]$ contains all tight elements of $[m_{j_1}, M_{j_1}]$. \square

Theorem 4.15. *For $t^* > 1$, Algorithm PUSHDOWN-REDUCE performs the exact same steps with inputs \mathcal{I} and \mathcal{I}' of Lemma 4.14 until iteration $t^* - 1$ when $\text{REDUCE}(j_2, t^* - 1, \mathcal{I}')$ is called. Hence compared to \mathcal{I} , the main loop of Algorithm PUSHDOWN-REDUCE terminates one step earlier with $t = t^* - 1$ when run with \mathcal{I}' .*

To prove Theorem 4.15 now we define elements that are no longer tight and elements that become tight in the new cover:

Lemma 4.16. *Let*

$$\begin{aligned} Z_1 &= \{x \text{ tight in } \mathcal{I} \text{ and } x \text{ not tight in } \mathcal{I}'\}, \\ Z_2 &= \{x \text{ not tight in } \mathcal{I} \text{ and } x \text{ tight in } \mathcal{I}'\}. \end{aligned}$$

Then

$$Z_1 \subseteq \{x : x \in [m_{j_2}, M_{j_1}], x \not\leq m_{j_1}\} \quad (4.4)$$

$$Z_2 \subseteq \{x : x \in [m_{j_1}, M_{j_1}], x \not\leq m_{j_2}\}. \quad (4.5)$$

Hence the same elements are tight in I_{j_1} for \mathcal{I} as in $[m_{j_2}, M_{j_1}]$ for \mathcal{I}' .

Proof. We get \mathcal{I}' from \mathcal{I} by removing $[m_{j_1}, M_{j_1}]$ and adding $[m_{j_2}, M_{j_1}]$ instead. Hence the elements of Z_1 should be contained in the latter but not in the former, and similarly the elements of Z_2 should be in the former but not in the latter interval. \square

Next we show that the algorithm proceeds identical for \mathcal{I} and \mathcal{I}' for $t < t^*$. The proof is based on the fact that the key elements used in defining $u_i^{(t)}$ do not belong to $Z_1 \cup Z_2$.

Lemma 4.17. *Let $u_i^{(t)}$ denote elements selected by Algorithm PUSHDOWN-REDUCE with input \mathcal{I}' with the convention that $u_{j_1}^{(t)}$ belongs to the modified interval $I'_{j_1} = [m_{j_2}, M_{j_1}]$. Then for all $t < t^*$, we have $u_i^{(t)} = u_i'^{(t)}$.*

Proof. By induction on $t \leq t^* - 1$, we will show $u_i^{(t)} = u_i'^{(t)}$. We prove the inductive hypothesis in three steps: we show for $i = 1, \dots, r$ that

- (i) $u_i^{(t)} \notin Z_1$;
- (ii) $u_i'^{(t)}$ exists; and
- (iii) $u_i'^{(t)} \notin Z_2$

The above three statements imply $u_i^{(t)} = u_i'^{(t)}$ as follows. For $t = 1$, the maximal tight elements are identical for $i \neq j_1$ by (i) and (iii), since $u_i'^{(1)}$ tight in \mathcal{I} implies $u_i^{(1)} \preceq u_i'^{(1)}$ and we have the opposite inequality when exchanging the role of the two elements. Also $u_{j_1}^{(1)} = u_{j_1}'^{(1)}$, since by Lemma 4.16, the tight elements of I_{j_1} in \mathcal{I} are the same as the tight elements of I'_{j_1} in \mathcal{I}' . For general t by induction on the step of defining $u_i'^{(t)}$, one can observe that element $u_i^{(t)}$ belongs to the set U of Procedure PUSHDOWN($i - 1, t, \mathcal{I}'$) and the same holds when exchanging the role of $u_i'^{(t)}$ and $u_i^{(t)}$. Thus the two elements must be equal.

Now we prove (i–iii). First of all for $i = j_1$ the tight elements of I_{j_1} in \mathcal{I} are the same as those of I'_{j_1} in \mathcal{I}' by Lemma 4.16, yielding (i–iii). Hence we assume $i \neq j_1$ next.

Proof of (i). Assume $u_i^{(t)} \in Z_1$. By Lemma 4.16, $m_{j_2} \preceq u_i^{(t)} \preceq M_{j_1}$ and $m_{j_1} \not\leq u_i^{(t)}$. Furthermore, since $m_{j_2} \preceq u_{j_1}^{(t^*)} \preceq u_{j_1}^{(t+1)} \preceq M_{j_1}$ we have $u_i^{(t)}$ and $u_{j_1}^{(t+1)}$ dependent. Using Corollary 4.10, $u_{j_1}^{(t+1)} \preceq u_i^{(t)}$, thus $m_{j_1} \preceq u_i^{(t)}$, a contradiction.

Proof of (ii). We show that $u_i'^{(t)}$ exists and $m_i \preceq u_i^{(t)} \preceq u_i'^{(t)}$. We proved above that $u_i^{(t)} \notin Z_1$ and hence $u_i^{(t)}$ remains tight in \mathcal{I}' . This immediately gives the result for $t = 1$. And for $t > 1$ we use the consequence of the inductive hypothesis that $u_h^{(t-1)} = u_h'^{(t-1)}$ for all h . This yields $u_i^{(t)} \in U$ for PUSHDOWN($i, t - 1, \mathcal{I}'$) that in turn implies that $u_i'^{(t)}$ exists and $u_i^{(t)} \preceq u_i'^{(t)}$.

Proof of (iii). Assume $u_i^{(t)} \in Z_2$. By Lemma 4.16, $m_{j_1} \preceq u_i^{(t)} \preceq M_{j_1}$, thus $u_i^{(t)}$ and $u_{j_1}^{(t+1)}$ are dependent. Observe furthermore $u_{j_1}^{(t+1)}$ is also tight in \mathcal{I}' . Hence by applying Lemma 4.3 for \mathcal{I}' , we get that either $u_{j_1}^{(t+1)} \preceq M_i$ or $m_i \preceq u_{j_1}^{(t+1)}$. In both cases we derive a contradiction with the definition of $u_{j_1}^{(t+1)}$ in Procedure PUSHDOWN(j_1, t, \mathcal{I}) by showing that certain $u_d^{(t)}$ may push $u_{j_1}^{(t+1)}$ down.

Case I: $u_{j_1}^{(t+1)} \preceq M_i$. By Lemma 4.16, we also get $m_{j_2} \not\preceq u_i^{(t)}$, which in turn implies $u_{j_1}^{(t+1)} \not\preceq u_i^{(t)}$, since $m_{j_2} \preceq u_{j_1}^{(t+1)}$. Because $u_{j_1}^{(t+1)}$ is tight in \mathcal{I}' and $u_{j_1}^{(t+1)} \preceq M_i$, we may apply Lemma 4.11 for \mathcal{I}' , $u_i^{(t)}$ and $z = u_{j_1}^{(t+1)}$. By the Lemma there exists $t_0 < t$ and $1 \leq d \leq r$ such that the element $u_d^{(t_0)}$ may push $u_{j_1}^{(t+1)}$ down. By induction $u_d^{(t_0)} = u_d^{(t_0)}$, and by Lemma 4.5, $u_d^{(t)}$ may also push $u_{j_1}^{(t+1)}$ down.

Case II: $m_i \preceq u_{j_1}^{(t+1)}$ and $u_{j_1}^{(t+1)} \not\preceq M_i$. As we have seen above, $m_i \preceq u_i^{(t)} \preceq u_i^{(t)}$. Thus $u_{j_1}^{(t+1)}$ and $u_i^{(t)}$ are dependent since their common lower and upper bounds are m_i and M_{j_1} , respectively. Hence in this case we have $d = i$: element $u_i^{(t)}$ may push $u_{j_1}^{(t+1)}$ down. The proof is complete. \square

We complete the proof of Theorem 4.15 by the following lemma.

Lemma 4.18. *When run with input \mathcal{I}' , Procedure REDUCE is called in iteration $t^* - 1$ with $j = j_2$.*

Proof. By Lemma 4.17, Procedure REDUCE cannot be called for \mathcal{I}' before iteration $t^* - 1$. Two things are left to prove: (i) in iteration $t^* - 1$, REDUCE($h, t^* - 1, \mathcal{I}'$) is not called for any $h < j_2$; and (ii) REDUCE($j_2, t^* - 1, \mathcal{I}'$) is called.

To prove (i), assume by contradiction that REDUCE($h, t^* - 1, \mathcal{I}'$) is called for some $h < j_2$, or equivalently, $U = \emptyset$ in PROCEDURE($h, t^* - 1, \mathcal{I}'$). We show that $u_h^{(t^*)} \in Z_1$. Indeed, by Lemma 4.17, $u_h^{(t^*-1)} = u_h^{(t^*-1)}$ for all h . Since no $u_h^{(t^*-1)}$ may push $u_{j_2}^{(t^*)}$ down, this yields that if $u_h^{(t^*)} \notin Z_1$, then $u_h^{(t^*)} \in U$, contradicting the assumption $U = \emptyset$.

By Lemma 4.16, $m_{j_2} \preceq u_h^{(t^*)} \preceq M_{j_1}$, thus q and $u_h^{(t^*)}$ are dependent. Element $u_h^{(t^*)}$ may not push q down, because it would contradict the fact that $\ell = j_2$ is minimal in a fixed ordering of the intervals so that $u_\ell^{(t^*)}$ may push q down. This means that $q \preceq M_h$. In addition, $q \not\preceq u_h^{(t^*)}$, since $m_{j_1} \preceq q$ and $m_{j_1} \not\preceq u_h^{(t^*)}$ by $u_h^{(t^*)} \in Z_1$. We can apply Lemma 4.11 for $u_h^{(t^*)}$ and $z = q$, which implies the existence of some $t_0 < t^*$ and $1 \leq d \leq r$ so that $u_d^{(t_0)}$ may push q down. By the second part of Lemma 4.12, $u_d^{(t_0)}$ may also push $u_{j_1}^{(t_0+1)}$ down, a contradiction.

For (ii), suppose for a contradiction that $u_{j_2}^{(t^*)}$ exists. Since $u_{j_2}^{(t^*)} \succeq m_{j_2}$, by Lemma 4.16, $u_{j_2}^{(t^*)} \notin Z_2$, hence $u_{j_2}^{(t^*)}$ is also tight in \mathcal{I} . We use again that by Lemma 4.17, $u_h^{(t^*-1)} = u_h^{(t^*-1)}$ for all h . This yields $u_{j_2}^{(t^*)} \in U$ for PUSHDOWN($j_2, t^* - 1, \mathcal{I}$), implying $u_{j_2}^{(t^*)} \preceq u_{j_2}^{(t^*)}$. By making use of Lemma 4.12, $u_{j_2}^{(t^*)} \preceq u_{j_2}^{(t^*)} \preceq M_{j_1}$.

We claim that $u_{j_2}^{(t^*)} \in Z_1$, contradicting the fact that $u_{j_2}^{(t^*)}$ is tight in \mathcal{I}' . As $m_{j_2} \preceq u_{j_2}^{(t^*)} \preceq M_{j_1}$ and $u_{j_2}^{(t^*)}$ is tight in \mathcal{I} , all we need to show is $m_{j_1} \not\preceq u_{j_2}^{(t^*)}$. Assume $m_{j_1} \preceq u_{j_2}^{(t^*)}$; this implies

$m_{j_1} \preceq u_{j_2}^{(t^*)} \preceq M_{j_1}$, thus $q \preceq u_{j_2}^{(t^*)}$ as q is the minimal tight element of $[m_{j_1}, M_{j_1}]$ in \mathcal{I} . In this case $u_{j_2}^{(t^*)}$ may not push q down, contradicting the selection of j_2 in Procedure REDUCE(j, t^*, \mathcal{I}). \square

4.2 Application for directed connectivity augmentation

In this section we give a reformulation of the above general algorithm which is applicable for the problem of directed node connectivity augmentation. The main difficulty is that we typically have an exponential size poset implicitly given as a set of (directed) cuts. We may either select an appropriate poset representation or implement the steps of the algorithm with direct reference to the underlying graph problem. We follow the second approach. We will show how all non trivial steps of the algorithm can be reduced to determining maximal tight elements in certain interval covers, which can be implemented as a sequence of BFS computations using some initial flow computations.

The key step in implementing Procedure PUSHDOWN for the underlying graph problems is the following reformulation of the main algorithm. We replace Procedure PUSHDOWN by an iterative method Procedure ALTERNATE-PUSHDOWN (see box) that selects a strictly descending sequence of tight elements $y_0 > y_1 > \dots > y_\ell$ with $y_0 = u_j^{(t)}$ and $y_\ell = u_j^{(t+1)}$ or terminates by Procedure REDUCE(j, t^*, \mathcal{I}). In the implementation for graph augmentation problems it is key to notice that in a single iteration of Procedure ALTERNATE-PUSHDOWN we only consider elements that may be pushed down by $u_i^{(t)}$ for a single value of i .

```

Procedure ALTERNATE-PUSHDOWN( $j, t, \mathcal{I}$ )
   $y_0 \leftarrow u_j^{(t)}$ ;  $h \leftarrow 0$ ;
  while exists  $i$  such that  $u_i^{(t)}$  may push  $y_h$  down do
     $U_h \leftarrow \{x : m_j \preceq x \preceq y_h, x \text{ tight and } u_i^{(t)} \text{ may not push } x \text{ down}\}$ 
    if  $U_h = \emptyset$  then
       $t^* \leftarrow t$ ;
      return REDUCE( $j, t^*, \mathcal{I}$ )
    else
       $y_{h+1} \leftarrow \text{maximal } x \in U_h$ ;
       $h \leftarrow h + 1$ 
  return  $y_h$ 

```

Lemma 4.19. *Procedures PUSHDOWN and ALTERNATE-PUSHDOWN return the same output.*

Proof. It follows straightforward from Lemma 4.6 that if $U_h \neq \emptyset$, then it has a unique maximal element, hence y_h for $h \geq 1$ is well defined.

If Procedure ALTERNATE-PUSHDOWN terminates by returning y_ℓ , then $y_\ell \in U$ for U as in Procedure PUSHDOWN. Thus $y_\ell \preceq u_j^{(t+1)}$. This shows that if Procedure PUSHDOWN terminates by calling Procedure REDUCE, then so does Procedure ALTERNATE-PUSHDOWN.

Consider now the case when $U \neq \emptyset$ in Procedure PUSHDOWN. We show that $y_h \succeq u_j^{(t+1)}$ for each $h \geq 0$. By contradiction, choose the smallest h with $y_h \not\succeq u_j^{(t+1)}$; thus $y_{h-1} \succeq y_h \vee u_j^{(t+1)} \succ y_h$. By the definition of U_{h-1} , $u_i^{(t)}$ may push $y_h \vee u_j^{(t+1)}$ down for some i . Using Lemma 4.6 again it may push either y_h or $u_j^{(t+1)}$ down, both leading to contradiction. Now we can conclude that if Procedure ALTERNATE-PUSHDOWN terminates by returning y_h , then both $y_h \preceq u_j^{(t+1)}$ and $y_h \succeq u_j^{(t+1)}$ hold, thus they are equal. \square

To compute y_h , consider the set of intervals $\mathcal{J}_{j,i} = \mathcal{I} - [m_i, M_i] + [m_i, M_j]$ with i as in Procedure ALTERNATE-PUSHDOWN. While $\mathcal{J}_{j,i}$ is not necessarily a cover of the entire poset, the following lemmas still hold:

Lemma 4.20. *All $x \in U_h$ are tight in $\mathcal{J}_{j,i}$.*

Proof. Notice x is either contained in both intervals $[m_i, M_i]$ and $[m_i, M_j]$ or in neither of them: if $m_i \preceq x$, then x and $u_i^{(t)}$ are dependent because m_i is a common lower and $u_i^{(t)} \vee y_h$ a common upper bound. Hence $x \preceq M_i$, since $u_i^{(t)}$ may not push x down. \square

Lemma 4.21. *Suppose $u_i^{(t)}$ may push y_h down. The set of intervals $\mathcal{J}_{j,i}$ covers all elements of I_j ; furthermore $y_{h+1} = y_h \wedge Q$, where Q is the maximal tight element of I_j in $\mathcal{J}_{j,i}$.*

Proof. For all $x \in I_j$, we have $x \in [m_i, M_j]$ if $x \in [m_i, M_i]$, hence the number of intervals covering x cannot be less in $\mathcal{J}_{j,i}$ than in \mathcal{I} , thus $\mathcal{J}_{j,i}$ covers all elements of I_j .

For the second part we first show that if I_j has any tight elements for $\mathcal{J}_{j,i}$, then there is a unique maximal among them. We cannot apply Lemma 4.1 directly since $\mathcal{J}_{j,i}$ is not a cover, but the claim holds for any $x, y \in I_i$, since $x, y, x \vee y$ and $x \wedge y$ are all covered by $\mathcal{J}_{j,i}$. Hence the existence of the unique maximal tight element follows. Since any element of I_i is covered in $\mathcal{J}_{j,i}$ by at least as many intervals as in \mathcal{I} , Q is also tight in \mathcal{I} .

Finally we let $z = y_h \wedge Q$ and show $z = y_{h+1}$. Notice that z is tight in \mathcal{I} as it is an intersection of two tight elements in \mathcal{I} . As $y_{h+1} \preceq y_h$ and y_h is tight in $\mathcal{J}_{j,i}$ by Lemma 4.20, we get $y_{h+1} \preceq Q$ and thus $y_{h+1} \preceq y_h \wedge Q = z$. For $z \preceq y_{h+1}$ we have to prove that $u_i^{(t)}$ may not push z down. Indeed, suppose that $u_i^{(t)}$ may push z down. Then $m_i \preceq z \not\preceq M_i$, hence by $z \preceq Q$ follows $Q \in [m_i, M_j]$. As Q is tight in $\mathcal{J}_{j,i}$, this implies that $Q \in [m_i, M_i]$, thus $z \preceq Q \preceq M_i$, a contradiction. \square

By the lemma, the basic step of Procedure ALTERNATE-PUSHDOWN consists of computing the maximum tight element of an interval for certain set of covering intervals. Furthermore, at the beginning of the algorithm $u_j^{(1)}$ is the maximum tight element of I_j . Now we turn our attention to the implementation of the steps of the algorithm for connectivity augmentation.

We use the reduction of node connectivity augmentation to poset covering as Claim 1.39: the minimal elements correspond to set pairs having a singleton tail and all the other nodes as head; maximal elements are found by exchanging the role of tails and heads. For each interval $I = [m_i, M_i] \in \mathcal{I}$ we augment the graph by an edge $s_i t_i$ with s_i corresponding to m_i and t_i corresponding to M_i as in the above reduction. If \mathcal{I} covers all poset elements in $[m_i, M_i]$, then the minimum s_i - t_i cut in the augmented graph has value at least k .

Algorithm PUSHDOWN-REDUCE(\mathcal{I}) will first be applied for a greedy cover \mathcal{I} (for example, including all possible intervals), and then subsequently for covers of decreasing cardinality, until we finally reach an optimal cover. We initialize PUSHDOWN-REDUCE(\mathcal{I}) by computing $|\mathcal{I}|$ maximum flows, one corresponding to each interval in \mathcal{I} . For interval $[m_j, M_j]$ we compute a maximum s_j - t_j flow. Since \mathcal{I} is a cover, the maximum flow value is at least k . If the s_j - t_j flow value is more than k , then $[m_j, M_j]$ contains no tight elements thus can be removed from the cover and the iteration PUSHDOWN-REDUCE(\mathcal{I}) is finished. Otherwise $u_j^{(1)}$ is the set pair corresponding to the value k cut with maximal tail that can be obtained by a breadth-first search from t_j on the graph obtained from the standard auxiliary graph in the Ford-Fulkerson algorithm by reverting the edges.

Lemma 4.22. *Consider the task of finding the maximum tight element of an interval $I_j = [m_j, M_j]$ for certain set of intervals $\mathcal{J}_{j,i}$ (as for example in Procedure ALTERNATE-PUSHDOWN) that cover I_j . Using the maximum s_j - t_j flow computed at the initialization for I_j , this step requires $O(1)$ breadth-first search (BFS) computations.*

Proof. Consider the maximum s_j - t_j flow computed at the initialization. We add an edge $s_i t_j$ to the graph and remove the edge $s_i t_i$. If the flow contains the removed edge, then we remove the single flow path containing it. We augment the resulting flow to a maximum flow by a single BFS computation. By another BFS starting from t_j we either obtain the maximum tight element or deduce that there are no tight elements and Procedure REDUCE can be called. \square

For implementing REDUCE, we need to find minimal tight elements of certain intervals and a sequence of changes in the interval cover by adding an interval and removing another. The first step can be performed by a BFS computation from the corresponding s_i ; for the second step we need to update the flows corresponding to the intervals $[m_j, M_j] \in \mathcal{I}$. For each $[m_j, M_j]$ in iteration s , we consider the maximum s_j - t_j flow, add an edge $s_{j_{s+1}} t_{j_s}$ to the graph and remove the edge $s_{j_s} t_{j_s}$. Again, if the flow contains the removed edge, then we remove the single flow path containing it, an augment the flow by a BFS computation.

4.2.1 Running times

To estimate the running time we need bounds for the number of intervals j and the length of a longest chain ℓ in the poset. At the initialization of PUSHDOWN-REDUCE we perform j max-flow

computations; then the dominating steps are finding elements y_h in Procedure ALTERNATE-PUSHDOWN. Since computing meets \vee and intersections \wedge of elements as well as checking whether $u \in I_i$ may push v down can be done in $O(1)$ time, this step is dominated by $O(1)$ BFS computations by Lemma 4.22.

Between two calls to Procedure REDUCE the total number of iterations in all calls to ALTERNATE-PUSHDOWN that compute certain y_h can be bounded by $j \cdot \ell$, since in each step we find a strictly smaller element of certain interval. This totals to $O(j \cdot \ell)$ BFS computations. For an iteration of REDUCE, we also have to do $O(j \cdot \ell)$ BFS computations. The total number of calls to Algorithm PUSHDOWN-REDUCE is bounded by j since the number of intervals decreases in each iteration. Hence we have $O(j^2)$ maximum flow and $O(j^2 \cdot \ell)$ BFS computations.

For the node-connectivity augmentation problem $\ell = O(n)$, and $j = O(n^2)$ since adding a complete digraph surely gives an $(n - 1)$ -connected digraph. Thus by the above estimations the running time is dominated by $O(n^5)$ BFS computations and $O(n^4)$ Max Flow Computations. As a BFS can be computed in time $O(n^2)$ and a Max Flow in time $O(n^3)$, the total running time can be bounded by $O(n^7)$.

4.3 Further remarks

While we have outlined only the implementation of the algorithm for directed connectivity augmentation, it can be done similarly for other applications, for example, ST -edge-connectivity augmentation. The existence of a strongly polynomial, or even polynomial combinatorial algorithm, however, remains open. This latter application demonstrates its importance as by ST -edge-connectivity we may have arbitrarily large connectivity requirement k .

One may wonder of how strong the generalizational power of the interval covering problem. Two algorithmically equivalent problems, Dilworth's chain cover and bipartite matching, are special cases of interval covers; our algorithm generalizes the standard augmenting path matching algorithm. One may ask whether the network flow problem as different algorithmic generalization of matchings could also fit into our framework. We might also hope that ideas such as capacity scaling, distance labeling and preflows [1] that give polynomial algorithms for network flows can be used in the construction of a polynomial algorithm for the interval covering problem.

Finally one may be interested in the efficiency of our algorithm for the particular problems that can be handled. Here particular implementations and good oracle choices are needed. We may want to reduce the number of mincut computations needed by polynomial size poset representations. One might also be able to give improvements in the sense of the Hopcroft–Karp matching algorithm [43].

Chapter 5

Local edge-connectivity augmentation

5.1 Coverings without partition constraints

5.1.1 From degree-prescription to augmentation

As indicated in Section 1.3, the augmentation Theorem 1.15 can easily be derived from the degree-prescribed Theorem 1.17. We include the argument here, since it is a starting point to similar deductions for the PCLECA problem. Only the SPSS-property of p is used and hence the deduction of Theorem 1.21 from Theorem 1.22 will be essentially the same.

Consider an arbitrary minimal vector $m' : V \rightarrow \mathbb{Z}_+$ satisfying (1.4). (That is, (1.4) gets violated if we decrease $m'(v)$ by one for any $v \in V$ with $m'(v) > 0$.) Let m be the result of the parity adjusting of m' . Theorem 1.15 follows from Theorem 1.17 by showing that for some subpartition \mathcal{X} of V , $m'(V) = p(\mathcal{X})$ and hence $m(V) = 2 \lceil \frac{1}{2}p(\mathcal{X}) \rceil$.

In this context, a set $X \subseteq V$ is called **tight** (with respect to m') if $m'(X) = p(X)$. A node $v \in V$ is **positive** if $m'(v) > 0$. The minimality of m' means that each positive v is contained in a tight set. Let \mathcal{X} be a collection of tight sets so that for every positive v , there exists an $X \in \mathcal{X}$ with $v \in X$. Choose \mathcal{X} with $\sum_{X \in \mathcal{X}} |X|$ minimal. We claim that \mathcal{X} is a subpartition of V . This completes the proof as it implies $m'(V) = p(\mathcal{X})$.

By the minimality, \mathcal{X} may not contain X and Y with $X \subseteq Y$. Assume $X, Y \in \mathcal{X}$ are intersecting. (1.7a) implies that $X \cap Y$ and $X \cup Y$ are also tight, while (1.7b) gives that $X - Y$ and $Y - X$ are tight and $m'(X \cap Y) = 0$. Let us replace X and Y by $X \cup Y$ in the first and by $X - Y$ and $Y - X$ in the second case; both contradict the minimal choice of \mathcal{X} .

5.1.2 Covering symmetric positively skew supermodular functions

We shall prove Theorem 1.41 in this section. We usually omit the index F and use $\nu = \nu_F$, $q = q_F$, $\mathcal{F} = \mathcal{F}_F$ etc. whenever clear from the context. The following is a well-known simple

property of the degree function¹.

Claim 5.1. *In a graph $G = (V, E)$, the degree function d satisfies the following for any $X, Y \subseteq V$:*

$$\begin{aligned} d(X) + d(Y) &= d(X \cap Y) + d(X \cup Y) + 2d(X, Y), \\ d(X) + d(Y) &= d(X - Y) + d(Y - X) + 2\bar{d}(X, Y) \end{aligned}$$

□

Together with the SPSS-property of p we get the following claim. (Recall the definition $q_F(X) = p(X) - d_F(X)$.)

Claim 5.2. *For any $X, Y \subseteq V$, with $p(X), p(Y) > 0$, at least one of the following inequalities hold:*

$$q(X) + q(Y) \leq q(X \cap Y) + q(X \cup Y) - 2d_F(X, Y), \quad (5.1a)$$

$$q(X) + q(Y) \leq q(X - Y) + q(Y - X) - 2\bar{d}_F(X, Y) \quad (5.1b)$$

□

When applying this claim, we usually omit checking $p(X), p(Y) > 0$, but this will always be easy to verify. An easy consequence is the following.

Claim 5.3. *If $q(X) = q(Y) = \nu$, then either $q(X \cap Y) = q(X \cup Y) = \nu$ or $q(X - Y) = q(Y - X) = \nu$. In addition, $d_F(X, Y) = 0$ in the first and $\bar{d}_F(X, Y) = 0$ in the second alternative. Consequently, \mathcal{F} is a subpartition of V .*

□

The next simple lemma describes the change in the values of q_F when a flipping is performed.

Lemma 5.4. *Consider a set $Z \subseteq V$. By flipping (xy, uv) , $q_F(Z)$ either remains unchanged or it increases or decreases by 2. It decreases by 2 if and only if both Z and $V - Z$ span exactly one of the two edges xy and uv . It increases by two if both Z and $V - Z$ span exactly one of the two edges xv and yu .*

□

We are now ready to prove Theorem 1.41. For a contradiction, assume $\nu \geq 2$. $|\mathcal{F}| \geq 2$ follows by the symmetry of p . Choose two sets $X, Y \in \mathcal{F}$, disjoint by Claim 5.3. (1.4) implies the existence of two edges $xy \in I_F(X)$, $uv \in I_F(Y)$ ($I_F(X)$ is the set of edges $xy \in F$ with $x, y \in F$). At this point, xy and uv are chosen arbitrarily; in the later part of the proof their choice we will be further specified.

Let F_1 and F_2 be the result of flipping (xy, uv) and (xy, vu) , respectively. We claim that either $F_1 \prec F$ or $F_2 \prec F$, leading to a contradiction.

¹Its directed counterpart is Claim 6.8.

Claim 5.5. *There exists no set Z with $q(Z) \geq \nu - 1$ crossing both X and Y .*

Proof. Assume for a contradiction that such a set exists. If (5.1a) held for X and Z , then $q(X \cap Z) + q(X \cup Z) \geq 2\nu - 1$. However, $q(X \cap Z) \leq \nu - 1$ by the minimal choice of X and hence $q(X \cup Z) = \nu$. Now Claim 5.3 yields a contradiction for $X \cup Z$ and Y . If (5.1b) held for X and Z , then similarly, $q(Z - X) = \nu$ and we get a contradiction for $Z - X$ and Y . \square

We call a set $Z \subseteq V$ **stable** if it does not contain a subset $U \subseteq Z$ with $q(U) = \nu$. The \preceq -minimal choice of F implies that either $\nu_{F_1} > \nu_F$ or $\nu_{F_1} = \nu_F$ and $|\mathcal{F}_{F_1}| \geq |\mathcal{F}_F|$. This enables us to derive an extremely useful structural property.

Lemma 5.6. *For $xy \in I_F(X)$, $uv \in I_F(Y)$, there exists a unique minimal stable $xv\bar{y}u$ -set² Z_{xv} and a unique minimal stable $\bar{x}v\bar{y}u$ -set Z_{yu} with $q(Z_{xv}) = q(Z_{yu}) = \nu - 2$, $Z_{xv} \cap Z_{yu} = \emptyset$. Furthermore, either (a) $q(Z_{xv} \cap X) = q(Z_{xv} \cup X) = \nu - 1$, $d_F(Z_{xv}, X) = 0$ or (b) $q(Z_{xv} - X) = q(X - Z_{xv}) = \nu - 1$, $\bar{d}_F(Z_{xv}, X) = 0$; analogous properties hold by changing the role of X and Y and also that of Z_{xv} and Z_{yu} .*

Proof. Lemma 5.4 and Claim 5.5 together imply $\nu_{F_1} \leq \nu_F$. Assume now $\nu_{F_1} = \nu_F = \nu$ but $|\mathcal{F}_{F_1}| \geq |\mathcal{F}_F|$. $X, Y \notin \mathcal{F}_{F_1}$, hence $|\mathcal{F}_{F_1} - \mathcal{F}_F| \geq 2$. This may only happen if there exist two disjoint stable sets Z_{xv} and Z_{yu} with $q(Z_{xv}) = q(Z_{yu}) = \nu - 2$, and Z_{xv} is an $xv\bar{y}u$ -set while Z_{yu} is a $\bar{x}v\bar{y}u$ -set. To see that a unique minimal Z_{xv} can be chosen, assume Z and Z' are two stable $xv\bar{y}u$ sets with $q(Z) = q(Z') = \nu - 2$. It suffices to show $q(Z \cap Z') = \nu - 2$. (5.1b) cannot hold for Z and Z' as it would give $q(Z - Z') = q(Z' - Z) = \nu$, contradicting the stability. Thus (5.1a) gives $q(Z \cap Z') + q(Z \cup Z') \geq 2\nu - 4$. Claim 5.5 implies that both terms are at most $\nu - 2$, hence $q(Z \cap Z') = \nu - 2$. The rest of the claim follows similarly, using Claim 5.2 for X and Z_{xv} . \square

The same argument for flipping (xy, vu) instead of (xy, uv) shows the existence of the sets Z_{xu}, Z_{yv} with analogous properties. This is an abuse of notation as the set Z_{xv} depends not only on the nodes x and v but on the edges xy and uv ; however, this should always be clear from the context. Claims 5.2 and 5.5 imply:

Claim 5.7. *At least one of the following alternatives hold:*

$$(a) \quad q(Z_{xu} \cap Z_{xv}) = q(Z_{xu} \cup Z_{xv}) = \nu - 1, \quad d_F(Z_{xu}, Z_{xv}) = 1, \quad Y \subseteq Z_{xu} \Delta Z_{xv};$$

$$(b) \quad q(Z_{xu} - Z_{xv}) = q(Z_{xv} - Z_{xu}) = \nu - 1, \quad \bar{d}_F(Z_{xu}, Z_{xv}) = 1, \quad (Z_{xu} \Delta Z_{xv}) \cap X = \emptyset.$$

There are analogue alternatives for Z_{yu} and Z_{yv} . \square

²By an $x\bar{y}$ -set we mean a set containing x and not containing y . We also use this notation for multiple nodes, for example, an $xv\bar{y}u$ -set contains x and v and does not contain y and u .

Lemma 5.8. *There exist subsets $X_0 \subseteq X$, $Y_0 \subseteq Y$ with $q_F(X_0) = q_F(Y_0) = \nu - 1$ and X_0, Y_0 minimal subject to these properties. Furthermore, if T is stable, $q(T) = \nu - 1$, $X_0 - T$ and $X_0 \cap T$ are nonempty, then $X_0 \cup T \supseteq X$. The same holds for X_0 and X replaced by Y_0 and Y .*

Proof. By Lemma 5.6, either $q(X \cap Z_{xv}) = \nu - 1$ or $q(X - Z_{xv}) = \nu - 1$, implying the existence of X_0 . For the second part, $T - X_0 \neq \emptyset$ by the minimality of X_0 ; (5.1b) cannot hold for X_0 and T since $q(X_0 - T) \leq \nu - 2$ also by the minimality and $q(T - X_0) \leq \nu - 1$ by the stability of T . Thus (5.1a) holds. Again, $q(X_0 \cap T) \leq \nu - 2$ by the minimality of X_0 and hence $q(X_0 \cup T) \geq \nu$ implying $X_0 \cup T \supseteq X$. \square

Since $\nu - 1 > 0$, (1.4) enables us to choose the edges xy, uv with the stronger property $xy \in I_F(X_0)$, $uv \in I_F(Y_0)$. Take alternative (b) in Claim 5.7. Then $Z_{xu} - Z_{xv}$ and $Z_{xv} - Z_{xu}$ fulfill the conditions on T in Lemma 5.8 for Y , giving that the nonempty set $Y - Y_0$ is contained in both, a contradiction as these sets are disjoint. Thus alternative (a) holds for Z_{xu}, Z_{xv} and similarly for Z_{yu}, Z_{yv} . Now $Z_{xu} \cap Z_{xv}$ and $Z_{yu} \cap Z_{yv}$ fulfill the conditions on T and hence both contain $X - X_0$, a contradiction again (they are disjoint as Z_{xu} and Z_{yv} have already been disjoint.) The proof of Theorem 1.41 is now complete.

5.1.3 New proof of Theorem 1.17

For $p(X) = (R(X) - d_G(X))^+$, we have the following slightly stronger version of Claim 5.2, with d_{G+F} instead of d_F :

Claim 5.9. *For any $X, Y \subseteq V$ with $p(X), p(Y) > 0$, at least one of the following inequalities hold:*

$$q(X) + q(Y) \leq q(X \cap Y) + q(X \cup Y) - 2d_{G+F}(X, Y), \quad (5.2a)$$

$$q(X) + q(Y) \leq q(X - Y) + q(Y - X) - 2\overline{d_{G+F}}(X, Y) \quad (5.2b)$$

Besides this, the only specific property of R we use is

$$R(X \cup Y) \leq \max\{R(X), R(Y)\} \text{ for any disjoint sets } X, Y \subseteq V, \quad (5.3)$$

straightforward from the definition of R .³ In fact, (5.3) will solely be used to prove Lemma 5.10.

To prove Theorem 1.17, choose a \preceq -minimal m -prescribed edge-set F ; $\nu_F \leq 1$ by Theorem 1.41. We are done if $\nu_F = 0$, therefore the only remaining case is $\nu_F = 1$.

Let us adapt the notation of the proof of Theorem 1.41. The argument of the proof fails for $\nu = 1$ since although X_0 and Y_0 exist, $I_F(X_0)$ or $I_F(Y_0)$ might be empty. Instead, we will use the following connectivity property:

³Actually, this property is valid for arbitrary (not necessarily disjoint) sets X and Y . In fact, if we require it for arbitrary sets, it will itself imply not only that R is skew-supermodular but also that it arises in the form (1.2) from a connectivity requirement function r . On the other hand, given a function R which is symmetric, skew-supermodular and satisfies (5.3), it does not follow that R arises in the form (1.2).

Lemma 5.10. *If $\nu = 1$, then there exists no $\emptyset \neq U \subsetneq X$ such that $d_G(U, X - U) = 0$ and $d_F(U, X - U) \leq 1$. The same holds for Y .*

Proof. Suppose, contrary to our claim, that such a set U existed. $R(X) \leq \max\{R(U), R(X-U)\}$ by (5.3). By symmetry, assume $R(X) \leq R(U)$. Also, $d_{G+F}(U) \leq d_{G+F}(X) - d_{G+F}(X - U, V - X) + 1$. By the minimal choice of X , $q(U) < q(X) = R(X) - d_{G+F}(X)$, implying $d_{G+F}(X - U, V - X) = 0$, hence $d_G(X - U) = 0$. $\nu = 1$ yields $R(X - U) \leq 1$, contradicting the assumption that there are no marginal sets. \square

In Claim 5.7, we can also write d_{G+F} instead of d_F because of the stronger Claim 5.9. Taking alternative (a), the disjoint sets $Z_{xu} - Z_{xv}$ and $Z_{xv} - Z_{xu}$ cover Y and the only edge connecting them is uv , a contradiction to Lemma 5.10. In alternative (b), xy is the only edge connecting $X \cap (Z_{xu} \cap Z_{xv})$ and $X - (Z_{xu} \cap Z_{xv})$, a contradiction again to Lemma 5.10. This completes the proof of Theorem 1.17.

5.1.4 New proof of Theorem 1.22

Assume now p is symmetric and positively crossing supermodular. Thus for crossing X, Y with $p(X), p(Y) > 0$, both (5.1a) and (5.1b), and also both alternatives in Lemma 5.6 and Claim 5.7 hold. We assume that (1.4) holds, but do not assume (1.6). Theorem 1.22 is an immediate consequence of the following:

Theorem 5.11. *Let F be a \preceq -minimal m -prescribed edge set. Either $\nu_F = 0$, or $\nu_F = 1$ and the following hold:*

- (i) \mathcal{F}_F forms a partition of V .
- (ii) $\dim(p) - 1 \geq |\mathcal{F}_F| + |F|$.
- (iii) There exists an edge set H covering p with $|H| = |\mathcal{F}_F| + |F|$.

We will need the following slight generalization of Lemma 1.20:

Lemma 5.12. *Let $\mathcal{P} = \{X_1, \dots, X_t\}$ be a subpartition of V so that $p(\bigcup_{i=1}^t X_i) = 0$, $p(X_1) = 1$ and $p(X_1 \cup X_j) > 0$ for any $j = 2, \dots, t$. Then \mathcal{P} is a p -full partition.*

Proof. Assume first \mathcal{P} is not a partition, that is, $V - \bigcup_{i=1}^t X_i \neq \emptyset$. By induction on $|\mathcal{I}|$, we prove that $p(\bigcup_{i \in \mathcal{I}} X_i) > 0$ for any $1 \in \mathcal{I} \subseteq \{1, \dots, t\}$. This will give a contradiction for $\mathcal{I} = \{1, \dots, t\}$. By the assumption, the claim is true for $|\mathcal{I}| \leq 2$. For some $z \in \mathcal{I} - \{1\}$, let $A = X_1 \cup X_z$ and $B = \bigcup_{i \in \mathcal{I} - z} X_i$. Now A and B are crossing and $p(A), p(B) > 0$, hence $p(A) + p(B) \leq p(A \cup B) + p(A \cap B)$. The claim follows as the LHS is at least 2, while $p(A \cap B) = 1$ and $A \cup B = \bigcup_{i \in \mathcal{I}} X_i$.

We have proved that $\bigcup_{i=1}^t X_i = V$. The same argument is still applicable for every $1 \in \mathcal{I} \subsetneq \{1, \dots, t\}$. Using the symmetry of p , we get that \mathcal{P} is p -full. \square

Proof of Theorem 5.11. $\nu_F \leq 1$ follows by Theorem 1.41; from now on, assume $\nu_F = 1$. Let us use the notation of the proof of Theorem 1.41: let $X, Y \in \mathcal{F}_F$, $xy \in I_F(X)$, $uv \in I_F(Y)$, Z_{xv} , Z_{yu} , Z_{xu} , Z_{yv} as in Lemma 5.6.

Lemma 5.13. (i) For each edge $xy \in I_F(X)$, there exist a unique maximal $\bar{x}y$ -set $D_{xy} \subseteq X$ and a unique maximal $x\bar{y}$ -set $D_{yx} \subseteq X$ with $q(D_{xy}) = q(D_{yx}) = 0$. Moreover, $D_{xy} \cap D_{yx} = \emptyset$, $D_{xy} \cup D_{yx} = X$. Analogous sets exist for edges in $I_F(Y)$.

(ii) For $xy \in I_F(X)$ and $uv \in I_F(Y)$, we have $Z_{xv} = D_{yx} \cup D_{uv}$.

(iii) For $xy \in I_F(X)$, the unique edge between D_{xy} and D_{yx} is xy . Furthermore, $d_F(X) = 0$.

(iv) For $xy, x'y' \in I_F(X)$, the sets D_{xy} and $D_{x'y'}$ are either disjoint or one contains the other or their union is X .

Proof. (i) For an arbitrary $uv \in I_F(Y)$, consider the set Z_{xv} . Both alternatives in Lemma 5.6 hold and thus $q(X \cap Z_{xv}) = q(X - Z_{xv}) = 0$. The existence of the unique maximal sets D_{xy} and D_{yx} easily follows by (5.1a). Also, (5.1a) would give a contradiction if $D_{xy} \cap D_{yx} \neq \emptyset$. $D_{xy} \cup D_{yx} = X$ follows by $X - Z_{xv} \subseteq D_{xy}$, $X \cap Z_{xv} \subseteq D_{yx}$. We have equality for both because of $D_{xy} \cap D_{yx} = \emptyset$.

(ii) By the above argument, we already have $Z_{xv} \cap X = D_{yx}$, $Z_{xv} \cap Y = D_{uv}$. Assume for a contradiction that $U = Z_{xv} - (X \cup Y) \neq \emptyset$. From Lemma 5.6, we obtain $q(Z_{xv} - X) = q(Z_{xv} - Y) = 0$. These two sets are crossing since $U \neq \emptyset$. (5.1a) gives $0 \leq q(Z_{xv}) + q(U)$ and thus $1 \leq q(U)$, a contradiction since Z_{xv} is stable.

(iii) Alternative (b) in Claim 5.7 gives the first part. The second part follows from (5.1b) applied for X and each of Z_{xv}, Z_{xu}, Z_{yv} and Z_{yu} for an arbitrary $uv \in I_F(Y)$.

(iv) can be derived easily using (5.1a) and (5.1b) for the sets $D_{xy}, D_{yx}, D_{x'y'}$ and $D_{y'y'}$. \square

These arguments work for all possible choices of X, Y, xy and uv . This enables us to derive the following nice structure. Let W_1, \dots, W_ℓ be the members of \mathcal{F}_F . Then each W_i admits a partition $\mathcal{W}_i = \{W_i^1, \dots, W_i^{s_i}\}$ satisfying the following:

- $d_F(W_i) = 0$ for $i = 1, \dots, \ell$.
- The edges in $I_F(W_i)$ are between different classes of \mathcal{W}_i , and $I_F(W_i)$ forms a spanning tree T_i if we contract the members of \mathcal{W}_i to single nodes.
- For an $uv \in I_F(W_i)$, the sets D_{uv} and D_{vu} are the unions of the members of \mathcal{W}_i corresponding to the connected components of $T_i - uv$ containing v and u , respectively.

Let $\mathcal{P} = \bigcup_{i=1}^{\ell} \mathcal{W}_i$. We claim that for some choice of $X_1 \in \mathcal{P}$, \mathcal{P} fulfils the conditions in Lemma 5.12. This immediately implies (i) and (ii) of the theorem. (iii) can be proved by induction: for some $i \neq j$, choose an arbitrary $x \in W_i$ and $v \in W_j$, increase $m(x)$ and $m(v)$ by

1 and add the edge xv to F . Clearly, if $|\mathcal{F}_F| > 2$ then it decreases by 1, and if $|\mathcal{F}_F| = 2$ then ν_F reduces to 0.

Let X_1 correspond to a leaf x in T_1 ; we may assume $X_1 = W_1^1 = D_{yx}$ for $xy \in I_F(W_1)$. Since $q(D_{yx}) = 0$ and $d_F(D_{yx}) = 1$, it follows that $p(X_1) = 1$. We need to prove $p(X_1 \cup W_i^j) > 0$ for any $1 \leq i \leq \ell$, $1 \leq j \leq s_i$.

First, consider the case $i > 1$. If W_i^j corresponds to a leaf in T_i , then $W_i^j = D_{uv}$ for some $uv \in I_F(W_i)$ and $X_1 \cup W_i^j = Z_{xv}$. We are done since $q(Z_{xv}) = -1$ and $d_F(Z_{xv}) = 2$. Next, assume that W_i^j is not a leaf. Let $uv \in I_F(W)$ be one of the edges entering W_i^j . Then $D_{uv} \supsetneq W_i^j$. Let $F' = \{u'v' \in I_F(D_{uv}), u' \in W_i^j - \{v\}\}$. Clearly, $D_{uv} = W_i^j \cup (\bigcup_{u'v' \in F'} D_{u'v'})$.

Let $A = D_{xy} \cup (\bigcup_{u'v' \in F'} D_{u'v'})$. This is the union of the sets $Z_{yv'}$ for $u'v' \in F'$. Recall that $p(D_{xy}) = p(Z_{yv'}) = 1$ for each $u'v'$. As in the proof of Lemma 5.12, the iterative application of (1.7a) for these sets gives $p(A) \geq 1$. Now (1.7b) for A and $Z_{xv} = D_{yx} \cup D_{uv}$ gives $2 \leq p(A) + p(Z_{xv}) \leq p(A - Z_{xv}) + p(Z_{xv} - A) = 1 + p(Z_{xv} - A)$, since $A - Z_{xv} = D_{xy}$. We are done since $Z_{xv} - A = X_1 \cup W_i^j$, the set we are interested in.

It remains to prove $p(X_1 \cup W_1^j) > 0$ for $2 \leq j \leq s_1$. Assume W_2^1 corresponds to a leaf in T_2 , $W_2^1 = D_{uv}$. Now $p(X_1 \cup W_2^1) = 1$, since $X_1 \cup W_2^1 = Z_{xv}$, and $p(W_1^j \cup W_2^1) \geq 1$ can be proved the same way as above. Then $2 \leq p(X_1 \cup W_2^1) + p(W_1^j \cup W_2^1) \leq p(W_2^1) + p(X_1 \cup W_1^j \cup W_2^1)$ and hence $p(B) \geq 1$ for $B = X_1 \cup W_1^j \cup W_2^1$. Note that $p(W_2) = 1$, since $q(W_2) = 1$, $d_F(W_2) = 0$. Applying (1.7b) for B and W_2 we get $2 \leq p(B) + p(W_2) \leq p(B - W_2) + p(W_2 - B) = p(X_1 \cup W_1^j) + p(D_{vu})$. We are done since $p(D_{vu}) = 1$. \square

5.2 Basic results on partition-constrained local edge-connectivity augmentation

5.2.1 Proof of Theorem 1.42

Let (F, φ) be an \vec{m} -prescribed legal edge set. For edges $xy, uv \in F$, the pair (xy, uv) is **flippable** if xy is an ij -edge and uv is an $i'j'$ -edge with $i \neq j'$, $j \neq i'$. In this case, flipping (xy, uv) with $\varphi'(xy, x) = i$, $\varphi'(xy, y) = j$, $\varphi'(uv, u) = i'$, $\varphi'(uv, v) = j'$ gives another \vec{m} -prescribed legal edge set (F', φ') . Notice that for two edges $xy, uv \in F$, at least one of (xy, uv) and (xy, vu) is a flippable pair.

Let us adapt the notation and results of Section 5.1 on covering SPSS-functions. Assume $\nu_F > 0$. The symmetry of p yields $|\mathcal{F}_F| \geq 2$. By way of contradiction, assume $|\mathcal{F}_F| \geq 3$. Let X, Y and W be three different (and thus disjoint) members of \mathcal{F}_F . By (1.4), there exist flippable edges $xy \in I_F(X)$, $uv \in I_F(Y)$.

If (xy, uv) is flippable, then Lemma 5.6 remains also valid in the current context. Lemma 5.8 is also applicable, as its proof used only the existence of a flippable edge pair and the SPSS-

property. We also need the following simple observation:

Claim 5.14. *There exists no set Z with $q(Z) \geq \nu - 2$ crossing all three sets X, Y and W .*

Proof. By Claim 5.2, $q(Z') \geq \nu - 1$ for either $Z' = Z \cup W$ or $Z' = Z - W$. This contradicts Claim 5.5. \square

A different argument is given for $\nu \geq 2$ and $\nu = 1$.

The case $\nu \geq 2$

For an edge $xy \in I_F(X)$, we say that the endnode y is **heavy** if there exists an $\bar{x}y$ -set $D \subseteq X$ with $q(D) = \nu - 1$. An endnode is **light** if it is not heavy. Heavy and light endnodes of edges in $I_F(Y)$ and in $I_F(W)$ can be defined in an analogous way.

Claim 5.15. *If y is a heavy endnode of the edge $xy \in I_F(X)$, then there exists a unique maximal $\bar{x}y$ -set $D_{xy} \subseteq X$ with $q(D_{xy}) = \nu - 1$. The analogous statement holds for edges in $I_F(Y)$ and in $I_F(W)$.*

Proof. Assume D and D' are two $\bar{x}y$ -sets with $D, D' \subseteq X$ and $q(D) = q(D') = \nu - 1$. We claim that $q(D \cup D') = \nu - 1$, implying the existence of a unique maximal D_{xy} . Indeed, if (5.2b) held for D and D' then $q(D - D') = q(D' - D) = \nu$ would follow, contradicting the fact that both are subsets of X . \square

Lemma 5.16. *For an edge $xy \in I_F(X)$, if the endnode x is light, then y is heavy. Furthermore, if x is light and (xy, uv) is flippable for some $uv \in I_F(Y)$, then v is a heavy endnode of uv . Also, $Z_{xv} \cap X = X - D_{xy}$ and $q(Z_{xv} - X) = \nu - 1$.*

Proof. Consider an edge $uv \in I_F(Y)$ with (xy, uv) flippable. Alternative (a) in Lemma 5.6 is excluded since x is light, hence $q(Z_{xv} - X) = q(X - Z_{xv}) = \nu - 1$. Now $D = X - Z_{xv}$ is an $\bar{x}y$ -set with $q(D) = \nu - 1$, implying that y is heavy. To see that v is also heavy, apply Claim 5.9 for $Z' = Z_{xv} - X$ and Y . (5.2b) cannot hold for Z' and Y . Indeed, $q(Z' - Y) \leq \nu - 1$ because Z' is stable, and $q(Y - Z') \leq \nu - 1$ by the minimality of Y . (5.2a) yields $q(Y \cap Z') = \nu - 1$ and hence v is heavy.

It is left to show that $Z_{xv} \cap X = X - D_{xy}$. On the one hand, $X - Z_{xv} \subseteq D_{xy}$ by the maximality of D_{xy} . On the other hand, assume that $Z_{xv} \cap D_{xy} \neq \emptyset$. (5.2a) cannot hold for Z_{xv} and D_{xy} as $d_F(Z_{xv}, D_{xy}) \geq 1$ and thus we would have $q(Z_{xv} \cap D_{xy}) + q(D_{xy} \cup Z_{xv}) \geq 2\nu - 1$, a contradiction. Hence (5.2b) applies, giving $q(Z_{xv} - D_{xy}) = \nu - 2$, contradicting the minimal choice of Z_{xv} . \square

Fix $X_0 \subseteq X$ be as in Lemma 5.8.

Lemma 5.17. *For every edge $xy \in I_F(X_0)$, exactly one of the two endnodes is heavy.*

Proof. According to the previous lemma, we only have to show that x and y cannot be both heavy. Indeed, assume D is an $x\bar{y}$ -set and D' is an $\bar{x}y$ -set with $q(D) = q(D') = \nu - 1$, $D, D' \subseteq X$, and both of them are chosen minimal to these properties. If D and D' are not disjoint, then they are crossing. Now (5.2b) would give that $D - D'$ and $D' - D$ are smaller sets with the same properties, while in the case of (5.2a), we have the contradictory $q(D \cup D') + q(D \cap D') \geq 2\nu$. However, the second part of Lemma 5.8 implies that $X - X_0$ is a subset of both D and D' , giving a contradiction. \square

Fix an $xy \in I_F(X_0)$ with heavy endnode y so that D_{xy} is maximal. Let $A = X - D_{xy}$. Again by Lemma 5.8, $A \subseteq X_0$, and $q(A) \leq \nu - 2$, since x is the light endnode of xy (and also by the minimality of X_0).

Claim 5.18. $I_F(A) = \emptyset$.

Proof. Indeed, assume that there exists an edge $x'y' \in I_F(A)$ with heavy endnode y' and consider the sets D_{xy} and $D_{x'y'}$. None of them is contained in the other because of $y' \notin D_{xy}$ and the maximality of D_{xy} . If (5.2b) held, then $q(D_{xy} - D_{x'y'}) = q(D_{x'y'} - D_{xy}) = \nu - 1$, a contradiction: by Lemma 5.8, both must be subsets of X_0 . In the case of (5.2a), we have $q(D_{xy} \cap D_{x'y'}) = q(D_{xy} \cup D_{x'y'}) = \nu - 1$, since $D_{xy} \cup D_{x'y'} \subseteq X - x'$. Now $D_{xy} \cup D_{x'y'}$ is a larger $\bar{x}'y'$ set, contradicting the maximality of $D_{x'y'}$. \square

Choose arbitrary edges $uv \in I_F(Y)$, $wz \in I_F(W)$ so that (xy, uv) and (xy, wz) are both flippable. Let $Z = Z_{xv}$ and $Z' = Z_{xz}$. Claim 5.14 implies $Z \cap W = Z' \cap Y = \emptyset$ and thus $Z - Z', Z' - Z \neq \emptyset$.

Lemma 5.19. xy is the only edge in $G + F$ incident to A .

Together with Claim 5.18, this will immediately lead to a contradiction. Indeed, $m(A) = 1$ because of $I_F(A) = \emptyset$. Now $d_G(A) = 0$ and (1.4) give $R(A) \leq 1$, hence A is a marginal set.

Proof. We already know by Lemma 5.16 that $Z \cap X = Z' \cap X = A$ and $q(Z - A) = q(Z' - A) = \nu - 1$. We shall prove $Z \cap Z' = A$. It suffices to verify that $Z \cap (Z' - A) = \emptyset$. Indeed, assume they intersected. If (5.2a) held for Z and $Z' - A$, then $q(Z \cap (Z' - A)) + q(Z \cup (Z' - A)) \geq 2\nu - 3$. This is a contradiction since the first term is at most $\nu - 1$ by the stability of Z , while the second is at most $\nu - 3$ by Claim 5.14. On the other hand, (5.2b) would give $q(Z - (Z' - A)) = \nu - 2$, a contradiction to the minimality of Z .

Hence A is the intersection of any two of the three sets X , Z and Z' . (5.2b) holds for any two of them, since (5.2a) is excluded by $q(A) \leq \nu - 2$ and $q(Z \cup Z') \leq \nu - 3$. (5.2b) gives $\overline{d_{G+F}}(Z, X) = \overline{d_{G+F}}(Z', X) = 0$, $\overline{d_{G+F}}(Z, Z') = 1$, leading to the desired conclusion. \square

The case $\nu = 1$

We will again use the connectivity property Lemma 5.10 as in the proof of Theorem 1.17. The next claim can be proved similarly.

Claim 5.20. *If $\nu = 1$ and Z is a stable set with $q(Z) = -1$, then Z is connected in $G + F$. \square*

Consider edges $xy \in I_F(X)$, $uv \in I_F(Y)$ and $wz \in I_F(W)$ so that (xy, uv) and (xy, wz) are flippable. Let us investigate the three sets X , $Z = Z_{xv}$ and $Z' = Z_{xz}$, pairwise crossing by Claim 5.14.

If (5.2b) held for X and Z , then $q(X - Z) = q(Z - X) = 0$, $\overline{d_{G+F}}(X, Z) = 0$. As in the proof of Lemma 5.19, it can be seen that Z' is disjoint from both $Z - X$ and $X - Z$. We get a contradiction to Claim 5.20, since $d_{G+F}(Z' \cap X, Z' - X) = 0$. Consequently, (5.2a) can be applied, giving $q(X \cup Z) = 0$. Let $A = X \cup Z$.

The same argument leads to $q(B) = 0$ for $B = X \cup Z_{yu}$. Assume now (5.2a) holds for A and B . $d_{G+F}(A, B) \geq 1$ because of the edge uv ; hence $q(A \cap B) = q(A \cup B) = 1$ and $d_{G+F}(A, B) = 1$ follows, giving $Y \subseteq A \cup B$. Since the sets Z_{xv} and Z_{yu} are disjoint, $A \cap B = X$ and thus $Y \subseteq A \Delta B$, giving a contradiction to Lemma 5.10 when applied for Y , as uv is a cut edge of Y .

On the other hand, (5.2b) for A and B gives $q(A - B) = q(B - A) = 0$, $\overline{d_{G+F}}(A, B) = 0$. Again, we can prove using the minimality of Z' and Claim 5.14 that Z' is disjoint from both $A - B$ and $B - A$, and we get a contradiction again to Claim 5.20 because of $d_{G+F}(Z' \cap X, Z' - X) = 0$.

5.2.2 Approximating with an additive error r_{\max}

In this section we shall prove Theorem 1.43. The key is the following simple corollary of Theorem 1.42. We say that a partition $\mathcal{Q} = (Q_1, \dots, Q_t)$ of V and a legal degree-prescription $\vec{m} = (m_1, \dots, m_t)$ are **compatible** if $m_i(v) = 0$ whenever $v \notin Q_i$.

Lemma 5.21. *Assume we have a partition $\mathcal{Q} = (Q_1, \dots, Q_t)$ of V with a compatible legal degree-prescription $\vec{m} = (m_1, \dots, m_t)$ satisfying (1.4). Let F be an \vec{m} -prescribed edge set. Then there exists a \mathcal{Q} -legal augmenting edge set H with $|H| = \frac{1}{2}m(V) + \nu_F$. Given \vec{m} , we can find such an H in polynomial time.*

Proof. We may assume that F is \preceq -minimal. The proof is by induction on ν_F . If $\nu_F = 0$ then $H = F$ is a \mathcal{Q} -legal augmenting edge set because of the compatibility. If $\nu_F > 0$, then $|\mathcal{F}_F| = 2$ by Theorem 1.42;. Let $\mathcal{F}_F = \{X, Y\}$. (1.4) yields two different colours i and j among two nodes $x \in X$, $v \in Y$ with $m_i(x), m_j(v) > 0$. Let us increase $m_i(x)$ and $m_j(v)$ by one; let \vec{m}' denote the resulting degree-prescription (which is clearly legal) and let $F' = F + xv$. Now $\nu_{F'} = \nu_F - 1$ and \vec{m}' is also compatible with \mathcal{Q} . Hence by induction, we have a \mathcal{Q} -legal augmenting edge set of size $\frac{1}{2}m'(V) + \nu_{F'} = \frac{1}{2}m(V) + \nu_F$, which is the desired conclusion.

From the algorithmic point of view, all we use from the extreme choice of F is that there is no improving flipping. This can be checked by a flow computation for each pair of edges of F , and the set system \mathcal{F}_F can also be determined via flow computations. \square

We do not estimate the running times as one can certainly gain a lot by careful implementations; this is beyond the scope of this chapter. Let us now turn to the proof of Theorem 1.43. We shall construct a legal degree-prescription \vec{m} compatible with \mathcal{Q} so that $m(V) = 2\Psi_{\mathcal{Q}}(G)$. Then the theorem will follow by the previous lemma, since r_{\max} is a trivial upper bound on ν_F .

First, let us choose a minimal m' satisfying (1.4) as in Section 5.1.1, regardless to the partition \mathcal{Q} . Let $m'_i(v) = m'(v)$ if $v \in Q_i$ and 0 otherwise. If (1.8) holds for m' , then we are done: consider the m we get from m' by parity adjusting. Clearly, $m(V) = \alpha(G)$.

Otherwise, there is exactly one j with $m'_j(V) > \frac{m'(V)}{2}$. We need the following simple claim (recall that a set X is called tight if $m'(X) = p(X)$ and $v \in V$ is positive if $m'(v) > 0$).

Claim 5.22. *If m' is minimal, then for each positive v there exists a unique minimal tight set X_v containing v . If $u \in X_v - v$, then the following m'' also satisfies (1.4): $m''(u) := m'(u) + 1$, $m''(v) := m'(v) - 1$, and $m''(z) := m'(z)$ otherwise.* \square

Consider now a positive $v \in Q_j$. If $X_v - Q_j \neq \emptyset$ then by the above claim, we can modify m' so that $m'_j(V)$ decreases by one. Let us iterate this procedure as long as possible. Either we arrive at an m' with $m'_j(V) = \frac{m'(V)}{2}$ and thus (1.8) is satisfied, or at a certain point, no more such modification is possible. Hence $m'_j(V) > \frac{m'(V)}{2}$ and $X_v \subseteq Q_j$ for every positive $v \in Q_j$. Using the uncrossing argument as in Section 5.1.1, we get a subpartition \mathcal{X} of Q_j with $p(\mathcal{X}) = m'_j(V)$. Afterwards, let us increase $m'(z)$ on an arbitrary node $z \in V - Q_j$ by $2m'_j(V) - m(V)$. The resulting m is a legal degree-prescription with $m(V) = \beta_j(G)$, as required.

5.2.3 Hydrae and medusae

For a partition \mathcal{H} , let $R_{\mathcal{H}} = \max_{Z \in \mathcal{H}} R(Z)$. Our aim is now to find a good characterization in order to decide whether a partition $\mathcal{H} = \{X^*, Y^*, C_1, \dots, C_{\ell}\}$ forms a hydra with heads X^* and Y^* . Let $\xi = \min\{R(X^*), R(Y^*)\}$, $\Xi = \max\{R(X^*), R(Y^*)\}$. Let $G_{\mathcal{H}}$ denote the graph on the node set $\{v_{X^*}, v_{Y^*}, v_{C_1}, \dots, v_{C_{\ell}}\}$ corresponding to the members of \mathcal{H} , and let $v_Z v_{Z'}$ be an edge if $R(Z, Z') \geq \xi$ for $Z, Z' \in \mathcal{H}$.

Theorem 5.23. *$\mathcal{H} = \{X^*, Y^*, C_1, \dots, C_{\ell}\}$ with $d_G(C_i, C_j) = 0$ for every $1 \leq i < j \leq \ell$ forms a hydra if and only if the following hold: $R_{\mathcal{H}} = \Xi$, and there is a path in $G_{\mathcal{H}}$ connecting v_{X^*} and v_{Y^*} . Furthermore, if $\xi < \Xi$ then there is a unique C_a with $R(C_a) = \Xi$, and $R(C_i, C_j) \leq \xi$ for every $1 \leq i < j \leq \ell$.*

Proof. Wlog. assume $R(X^*) = \Xi$, $R(Y^*) = \xi$. Let us show the necessity of the conditions first. $R_{\mathcal{H}} > \Xi$ means that for some i, j , $R_{\mathcal{H}} = R(C_i, C_j) > \Xi$. Now (1.5a) cannot hold for $X^* \cup C_i$ and

$X^* \cup C_j$. Next, assume there is no path in $G_{\mathcal{H}}$ between v_{X^*} and v_{Y^*} . Let I denote the set of those indices i for which v_{C_i} can be reached from v_{X^*} , and let $J = \{1, \dots, \ell\} - I$. Then (1.5a) cannot hold with equality for $Z = X^* \cup (\bigcup_{i \in I} C_i)$ and $Z' = X^* \cup (\bigcup_{j \in J} C_j)$, since $R(Z) < \xi$, $R(Z') = \Xi$, but $R(Z \cap Z') = \Xi$ and $R(Z \cup Z') = \xi$.

In the case of $\xi < \Xi$, assume first $R(C_i, C_j) > \xi$ for some $i \neq j$. Now (1.5a) cannot hold for $Z = Y^* \cup C_i$ and $Z' = Y^* \cup C_j$. Indeed, it is easy to see that $R(Z \cup Z') \leq \max\{R(Z), R(Z')\}$, and $\min\{R(Z), R(Z')\} > R(Z \cap Z')$. Assume next that there are multiple indices i with $R(X^*, C_i) = \Xi$. Let I and J be the partition of such indices into two nonempty sets. For $Z = X^* \cup (\bigcup_{i \in I} C_i)$ and $Z' = X^* \cup (\bigcup_{j \in J} C_j)$, we get a contradiction since $R(Z) = R(Z') = R(Z \cap Z') = \Xi$, although $R(Z \cup Z') < \Xi$.

Sufficiency is straightforward if $\Xi = \xi$ since the path in $G_{\mathcal{H}}$ between v_{X^*} and v_{Y^*} guarantees $R(X^* \cup (\bigcup_{i \in I} C_i)) = R(Y^* \cup (\bigcup_{i \in I} C_i)) = \Xi$ for arbitrary $I \subseteq \{1, \dots, \ell\}$. It is also easy to verify the definition for $\xi < \Xi$ using the path in $G_{\mathcal{H}}$ and the uniqueness of C_a . This is left to the reader. \square

In the rest of this section, we list some useful properties of hydrae, needed for proving the $\max \leq \min$ direction of the conjectures and Theorem 5.30. $\mathcal{H} = \{X^*, Y^*, C_1, \dots, C_\ell\}$ will always denote a hydra with heads X^* and Y^* . The following two lemmas can be proved by a simple induction based on the properties in Definition 1.44.

Lemma 5.24. *For a subset $I \subseteq \{1, \dots, \ell\}$,*

$$p(X^* \cup (\bigcup_{i \in I} C_i)) - p(X^*) = \sum_{i \in I} (p(X^* \cup C_i) - p(X^*)),$$

and the same holds for X^ substituted by Y^* .* \square

Let us fix a colour h . We say that an edge $xy \in V^2$ is a **ordinary edge** w.r.t. \mathcal{H} and h , if $m_h(x) > 0$ and $m_i(y) > 0$, for some $i \neq h$ and furthermore, x and y are in one of the following three configurations: (a) $x \in X^*$, $y \in Y^*$; (b) $x \in Y^*$, $y \in X^*$; or (c) $x \in \bigcup C_i$ and $y \in X^* \cup Y^*$.

Lemma 5.25. *(i) Let $xy \in V^2$ be a ordinary edge. Consider the graph $G' = G + xy$ and the degree-prescription \vec{m}' with $m'_h(x) = m_h(x) - 1$, $m'_i(y) = m_i(y) - 1$ and $m'_j(z) = m_j(z)$ otherwise. A tentacle C_i is h -odd for G', \vec{m}', p' if and only if it is h -odd for G, \vec{m}, p .⁴*

(ii) $\mathcal{H}' = \{X^, Y^*, C_1 \cup C_2, C_3, \dots, C_\ell\}$ is also a hydra.*

(iii) $\mathcal{H}' = \{X^ \cup C_1, Y^*, C_2, \dots, C_\ell\}$ is also a hydra. Moreover, a tentacle C_i is h -odd in \mathcal{H}' if and only if it is h -odd in \mathcal{H} .* \square

⁴Note that p is also dependent from G .

Unlike the previous two, the next lemma is not a direct consequence of the definition, however, follows easily from the structural characterization, Theorem 5.23.

Lemma 5.26. *For any tentacle C_i , $p(C_i \cup X^*) + p(C_i \cup Y^*) = p(X^*) + p(Y^*)$. \square*

An important consequence of this lemma is that C_i is h -odd if and only if $p(C_i \cup Y^*) - p(Y^*) + m_h(C_i)$ is odd, that is, X^* can be replaced by Y^* in the definition of h -odd tentacles.

In the next definition we define the subclass of hydrae, which plays a central role in the proof of Theorem 5.30.

Definition 5.27. The partition $\mathcal{H} = \{X^*, Y^*, C_1, \dots, C_\ell\}$ forms a **medusa** in G with **heads** X^*, Y^* and **tentacles** C_i if

- (i) $d_G(C_i, C_j) = 0$ for every $1 \leq i < j \leq \ell$; and
- (ii) $R(C_i) < \xi = \min\{R(X^*), R(Y^*)\}$ for at least $\ell - 1$ different values of $i \in \{1, \dots, \ell\}$.

Theorem 5.23 immediately implies that all medusae are hydrae. Indeed, if $R(C_i) < \xi$ holds for every tentacle C_i , then $R(X^*, Y^*) = \Xi = \xi$. If there is a single exceptional tentacle C_a , then either $G_{\mathcal{H}}$ contains the edge $v_{X^*}v_{Y^*}$ or the path $v_{X^*}v_{C_a}v_{Y^*}$. Notice that the underlying partition of a C_4 -configuration forms a hydra, however, not a medusa.

We give another, equivalent characterization of medusae. Let $\mathcal{H} = \{X^*, Y^*, C_1, \dots, C_\ell\}$ be a partition of V . For $1 \leq i, j \leq \ell$, $i \neq j$, we say that Z, Z' is a **separating pair** for i and j if both sets are unions of some components of \mathcal{H} ; furthermore, $C_i \subseteq Z \cap Z'$, $C_j \cap (Z \cup Z') = \emptyset$, $X^* \subseteq Z - Z'$ and $Y^* \subseteq Z' - Z$. For $1 \leq t \leq \ell$, we say that the separating pair Z, Z' is **coherent with t** if either $C_t \subseteq (Z \cap Z')$ or $C_t \cap (Z \cup Z') = \emptyset$. (Note that Z and Z' is always coherent with i and j .)

Lemma 5.28. *Let $\mathcal{H} = \{X^*, Y^*, C_1, \dots, C_\ell\}$ be a partition with $d_G(C_i, C_j) = 0$ for every $1 \leq i < j \leq \ell$. \mathcal{H} forms a medusa with heads X^* and Y^* if and only if for any $1 \leq i, j, t \leq \ell$, $i \neq j$, there exists a separating pair Z, Z' for i and j coherent with t , so that (1.5a) does not hold for Z and Z' .*

Proof. If \mathcal{H} is a medusa, then $Z = X^* \cup C_i$ and $Z' = Y^* \cup C_i$ is a separating pair for i and any $j \neq i$, coherent with t for any $1 \leq t \leq \ell$. For the other direction, let us use Ξ and ξ as before. We shall first prove $R(C_i, C_j) < \Xi$ for any $i \neq j$. By way of contradiction, assume $R_{\mathcal{H}} = R(C_i, C_j)$ for some $i \neq j$. Then (1.5a) clearly holds for any pair Z, Z' separating i and j . We also get a contradiction if there existed $i \neq j$ with $R(C_i) = R(C_j) = \Xi$ (and thus $R(C_i, X^* \cup Y^*) = R(C_j, X^* \cup Y^*) = \Xi$). In the case of $\xi = \Xi$ it already follows that \mathcal{H} is a medusa.

If $\xi < \Xi$ then wlog. assume $\Xi = R(X^*, C_a)$. By the argument above, there is a unique such a . Let ξ' be the second largest connectivity value between different classes of \mathcal{H} ($\xi \leq \xi'$). It suffices to prove that ξ' may occur only between Y^* and X^* or between Y^* and C_a .

Indeed, if $\xi' = R(C_i, C_j)$, we show that (1.5a) holds for any pair Z, Z' separating i and j , coherent with a . If $C_a \subseteq Z \cap Z'$, then $R(Z) + R(Z') = \Xi + \xi'$, $R(Z \cap Z') = \Xi$ and $R(Z \cup Z') = \xi'$. If $C_a \cap (Z \cup Z') = \emptyset$, then $R(Z) = R(Z \cup Z') = \Xi$ and $R(Z') = R(Z \cap Z') = \xi'$. Finally, if $\xi' = R(C_i, X^* \cup Y^*)$ for $i \neq a$, then we get a contradiction for any pair Z, Z' separating i and a . \square

5.2.4 $\max \leq \min$ in Conjectures 1.45 and 1.46

$\max \leq \min$ in Conjecture 1.45 is established by the following lemma:

Lemma 5.29. *Let us be given a hydra $\mathcal{H} = \{X^*, Y^*, C_1, \dots, C_\ell\}$, a legal degree-prescription $\vec{m} = (m_1, \dots, m_t)$, a fixed $1 \leq h \leq t$, and an arbitrary \vec{m} -prescribed legal edge set (F, φ) . Then $\nu_F \geq \tau_h(G, r, \vec{m}, \mathcal{H})$.*

Note that (1.4) is not being assumed.

Proof. The proof is by induction on $m(V)$. First, we shall prove that if $\vec{m} \equiv 0$, then the maximum value of p is at least $\tau_h(G, r, 0, \mathcal{H})$. (This maximum value equals ν_F for $F = \emptyset$, the unique \vec{m} -prescribed legal edge-set.) For an h -odd tentacle C_i , $p(X^* \cup C_i) - p(X^*)$ is odd. Let $I = \{i : p(X^* \cup C_i) - p(X^*) > 0\}$ and $J = \{1, \dots, \ell\} - I$. By Lemma 5.26, $p(Y^* \cup C_j) - p(Y^*) \geq 0$ for every $j \in J$. Furthermore, if C_j is h -odd, then we have strong inequality here. The number of such indices is at least $\chi_h - |I|$. Let $X = X^* \cup (\bigcup_{i \in I} C_i)$ and $Y = Y^* \cup (\bigcup_{i \in J} C_j)$. By Lemma 5.24, $p(X) - p(X^*) \geq |I|$ and $p(Y) - p(Y^*) \geq \chi_h - |I|$. $Y = V - X$ and thus $p(X) = p(Y)$. Therefore

$$\tau_h(G, r, \vec{m}, \mathcal{H}) = \frac{1}{2} (\chi_h + p(X^*) + p(Y^*)) \leq \frac{1}{2} (p(X) + p(Y)) = p(X) \leq \nu_\emptyset,$$

proving the claim.

Next, assume $\vec{m} \not\equiv \emptyset$, and let $uv \in F$ be an arbitrary edge. Let $a = \varphi(uv, u)$ and $b = \varphi(uv, v)$. We apply induction for $G' = G + uv$, $F' = F - uv$ and \vec{m}' , where \vec{m}' arises from \vec{m} by decreasing $m_a(u)$ and $m_b(v)$ by one. Let $\mathcal{H}' = \mathcal{H}$ unless u and v lie in different tentacles. If $u \in C_i$, $v \in C_j$ for $i \neq j$, then let us replace the tentacles C_i and C_j by $C_i \cup C_j$. We shall prove $\tau_h(G, r, \vec{m}, \mathcal{H}) \leq \tau_h(G', r, \vec{m}', \mathcal{H}')$. implying the claim.

It is a routine to check this for any possible configuration of uv and \mathcal{H} . For example, if uv is a ordinary edge w.r.t. to \mathcal{H} and \vec{m} , then we may apply Lemma 5.25(i). Let us now analyze the least trivial case when $u \in C_i$, $v \in C_j$ for $i \neq j$, and both C_i and C_j are h -odd in \mathcal{H} . If $h \in \{a, b\}$ then $C_i \cup C_j$ is h -odd in \mathcal{H}' , hence χ_h decreases by 1. However, the term $m_h(\bigcup C_i) - m(V)$ increases by one; all other terms are left unchanged. On the other hand, if $h \notin \{a, b\}$, then

$C_i \cup C_j$ is not h -odd in \mathcal{H}' , thus χ_h decreases by two, but $m_h(\bigcup C_i) - m(V)$ increases also by two. We leave it to the reader to verify the remaining cases. \square

We will also use this lemma for the $\max \leq \min$ direction of Conjecture 1.46. Let \mathcal{F} be an arbitrary legal augmenting edge set. Since \mathcal{Z} is an h -subpartition, F contains at least $p(\mathcal{Z})$ edges incident to the classes of \mathcal{Z} . Let $F_1 \subseteq F$ be an arbitrary subset of such edges with $|F_1| = p(\mathcal{Z})$; let $F_2 = F - F_1$. Clearly, $\nu_{F_1} \leq |F_2|$.

Let us define \vec{m} as follows. Let $m(v) = d_{F_1}(v)$ for $v \in V$, and let $m_i(v) = m(v)$ if $v \in Q_i$ and $m_i(v) = 0$ otherwise. In particular, $\sum(p(\mathcal{Z}) : Z \in \mathcal{Z}, Z \subseteq C_i) = m_h(C_i)$ for arbitrary tentacle C_i , hence C_i is h -odd if and only if it is h -toxic. Also, $p(\mathcal{Z}) = m_h(\bigcup C_i)$. These observations yield

$$\tau'_h(G, r, \mathcal{H}, \mathcal{Z}) = \tau_h(G, r, \vec{m}, \mathcal{H}) + \frac{1}{2}m(V).$$

Since $\frac{1}{2}m(V) = |F_1|$, by Lemma 5.29 we obtain

$$\tau'_h(G, r, \mathcal{H}, \mathcal{Z}) \leq \nu_{F_1} + |F_1| \leq |F_2| + |F_1| = |F|.$$

5.3 Towards proving the conjectures

In this section, we shall prove Conjecture 1.45 in a special setting.

Theorem 5.30. *Let (F, φ) be a \preceq -minimal $\vec{m} = (m_1, m_2)$ -prescribed legal edge set as in Conjecture 1.45. If $\nu_F \geq 2$ and $\bigcup \mathcal{F}_F = V$, then $\nu_F = \tau(G, r, \vec{m})$. Moreover, there is a medusa \mathcal{H} giving the optimum value.*

As we have already seen (e.g. in Section 5.2.1), the cases $\nu = 1$ and $\nu \geq 2$ are of different nature. We investigate here only the case $\nu \geq 2$. We already know $|\mathcal{F}_F| = 2$ by Theorem 1.42. As before, let X and Y denote its two members. Hence the assumption of the theorem is $X \cup Y = V$. An important consequence is that $q(Z) = \nu$ implies $Z = X$ or $Z = Y$.

The proof relies on the results of Section 5.2.1. So far, the only way of using the extreme choice of F has been that no improving flipping exists. Another operation will also be needed here. By **hexa-flipping** (xy, uv, wz) for three 12-edges $xy, uv, wz \in F$, we mean replacing F by $F' = F - \{xy, uv, wz\} + \{xv, uz, wy\}$, where the new edges are defined as 12-edges. Actually, this is a sequence of two flippings: flipping xy and uv first, then flipping uy and zw , yet it is easier to handle these two flippings together. The next simple lemma describes the changes in the values of q_F by a hexa-flipping.

Lemma 5.31. *Consider a set $Z \subseteq V$. By hexa-flipping (xy, uv, wz) , $q_F(Z)$ either remains unchanged or it increases or decreases by 2. It increases by 2 if and only if Z intersects the set $\{x, y, u, v, w, z\}$ in one of the following six sets or in the complement of one: $\{x, v\}$, $\{u, z\}$, $\{w, y\}$, $\{x, v, w\}$, $\{x, v, z\}$, $\{u, z, x\}$. \square*

We do not formulate the analogous characterization for the sets with $q_F(Z)$ decreasing since we will not need it. Let us call a set Z with $q_{F'}(Z) = q_F(Z) + 2$ an **increasing set** (w.r.t. the hexa-flipping).

Consider the minimal sets $X_0 \subseteq X$, $Y_0 \subseteq Y$ with $q(X_0) = q(Y_0) = \nu - 1$ as in Lemma 5.8⁵, and choose 12-edges $xy \in I_F(X_0)$, $uv \in I_F(Y_0)$. By Lemma 5.17, exactly one of the two endnodes of xy is light; wlog. assume this is x . By Lemma 5.16, the 2-coloured endnode of each edge in $I_F(Y)$ is heavy. This holds in particular for uv , and by changing the role of X and Y we can conclude that the 2-coloured endnode of all edges in $I_F(X) \cup I_F(Y)$ is heavy.

Our aim is to construct a hydra \mathcal{H} with $\tau_1(G, r, \vec{m}, \mathcal{H}) = \nu_F$. For this, further investigation of the structure of the edge set F is needed. We start by formulating a sequence of lemmas which together provide the construction; the proofs are postponed.

First, we extend the results of Section 5.2.1 and prove, in particular, that the 1-coloured endnode of all edges in $I_F(X) \cup I_F(Y)$ is light. For every 12-edge $xy \in I_F(X) \cup I_F(Y)$, consider the $\bar{x}y$ -set D_{xy} as in Claim 5.15. Let $A_{xy} = X - D_{xy}$ if $xy \in I_F(X)$, and $A_{xy} = Y - D_{xy}$ if $xy \in I_F(Y)$. Recall that a set $Z \subseteq V$ has been called stable if there exists no $U \subseteq Z$ with $q(U) = \nu$. Accordingly, we call a set $Z \subseteq V$ **steady**, if it has no subset U with $q(U) \geq \nu - 1$. In the next lemma, we prove, among other structural properties, that all sets A_{xy} are steady (in fact, we assert a slightly stronger property).

Lemma 5.32. *(i) Let $xy \in I_F(X)$ be an arbitrary 12-edge. Then x is a light and y a heavy endnode of xy . The set A_{xy} is steady, moreover, there exists no set $Z \subseteq V$ with $q(Z) = \nu - 1$, $y \notin Z$ and $Z \cap A_{xy} \neq \emptyset$. For an arbitrary 12-edge $wz \in I_F(Y)$, we have $Z_{wz} \cap X = A_{xy}$ and $q(D_{wz} \cup A_{xy}) = \nu - 2$.*

(ii) If $wz \in I_F(A_{xy})$ is an 12-edge, then $A_{wz} \subseteq A_{xy}$, $q(D_{wz} \cup A_{wz}) = \nu - 2$ and $d_{G+F}(A_{wz}, A_{xy} - A_{wz}) = 1$.

(iii) For 12-edges $xy, wz \in I_F(X)$ we have $d_{G+F}(A_{xy}, A_{wz}) = 0$.

(iv) If $wz \in F$ is an 12-edge with $z \in A_{xy}$, then $w \in A_{xy}$.

Analogous statements hold when exchanging the role of X and Y .

If the set systems $\mathcal{A}^0 = \{A_{xy} : xy \in I_F(X)\}$ and $\mathcal{B}^0 = \{A_{xy} : xy \in I_F(Y)\}$ were laminar, then we would already be ready to construct an optimal hydra \mathcal{H} . Unfortunately, this is not necessarily true, and thus these set systems are needed to be uncrossed. The uncrossing has to be done very carefully as we shall keep the valuable structural properties asserted in the previous lemma. This motivates the following definitions.

Assume $U, T \subseteq X$ are steady sets with $q(X - T) = q(X - U) = \nu - 1$ and $T \subsetneq U$. We say that T is a **descendant of U** if $q(T \cup (X - U)) = \nu - 2$, $d_{G+F}(T, U - T) = 1$ and there is a

⁵Recall from Section 5.2.1 that this lemma is also valid in the context of the PCLECA problem.

(unique) 12-edge $wz \in F$ from T to $U - T$. For example, if $xy, wz \in I_F(X)$ are 12-edges with $wz \in I_F(A_{xy})$, then Lemma 5.32(ii) states that A_{wz} is a descendant of A_{xy} .

We say that a set system \mathcal{A}' **blocks** the 12-edge $xy \in I_F(X)$ if \mathcal{A}' contains an $x\bar{y}$ -set. Analogously, \mathcal{A}' **blocks** the edge set $F' \subseteq I_F(X)$ if it blocks each edge in F' .

Definition 5.33. For a set $F' \subseteq I_F(X)$ of 12-edges and a set system \mathcal{A}' of subsets of X , we say that \mathcal{A}' is a **witness system for F'** if the following hold.

- (a) \mathcal{A}' is laminar, and for every $A \in \mathcal{A}'$, A is a steady set, $q(X - A) = \nu - 1$, $d_F(A, X - A) > 0$.
- (b) \mathcal{A}' blocks F' .
- (c) For each non-maximal $A \in \mathcal{A}'$, let $U \in \mathcal{A}' - A$ be the smallest set containing A . Then A is a descendant of U .

Descendants and witness systems for subsets of $I_F(Y)$ can be defined analogously. Let \mathcal{A}' and \mathcal{B}' be arbitrary sets of subsets of X and Y , respectively. \mathcal{A}' and \mathcal{B}' are called **linked** if

$$q(A \cup (Y - B)) = \nu - 2 \text{ holds for every } A \in \mathcal{A}, B \in \mathcal{B}. \quad (5.4)$$

Lemma 5.34. *There exist linked witness systems \mathcal{A}, \mathcal{B} for $I_F(X)$ and $I_F(Y)$, respectively.*




Figure 5.1: Illustration of Lemma 5.34. The sets in X form a witness system \mathcal{A} and those in Y form \mathcal{B} . The 1- and 2-endnodes of edges in F are denoted by circles and rectangles, respectively.

Consider now the witness systems \mathcal{A} and \mathcal{B} as in the previous lemma. Let \mathcal{C}^X denote the underlying subpartition of \mathcal{A} , that is, \mathcal{C}^X contains the minimal members of \mathcal{A} , and for each non-minimal member $A \in \mathcal{A}$, let \mathcal{C}^X contain $C = A - \bigcup\{A' : A' \in \mathcal{A}, A' \subsetneq A\}$. Let us say that A is the **corresponding member for** C in \mathcal{A} . Let $T_X = \bigcup \mathcal{C}^X = \bigcup \mathcal{A}$, and $X^* = X - T_X$. Define \mathcal{C}^Y , T^Y and Y^* the analogous way from \mathcal{B} . Let $\mathcal{C} = \mathcal{C}^X \cup \mathcal{C}^Y$, and let us denote its members by $\mathcal{C} = \{C_1, \dots, C_\ell\}$. The next lemma completes the proof of Theorem 5.30.

Lemma 5.35. *The partition $\mathcal{H} = \{X^*, Y^*, C_1, \dots, C_\ell\}$ forms a medusa with heads X^* and Y^* and tentacles C_i . Moreover, $\tau_1(G, r, \vec{m}, \mathcal{H}) = \nu_F$.*

5.3.1 Proofs of the Lemmas

Proof of Lemma 5.32. (i) It is enough to prove that A_{xy} is a steady set. Indeed, if Z were a set as in the conditions, then $Z \cap Y = \emptyset$ by Claim 5.5. Assume $Z \cap D_{xy} \neq \emptyset$; we claim that (5.2a) cannot hold for Z and D_{xy} . Indeed, from $x \in Z$ we would obtain $q(Z \cup D_{xy}) + q(Z \cap D_{xy}) \geq 2\nu$, while if $x \notin Z$ then $q(Z \cup D_{xy}) \geq \nu - 1$ gives a contradiction to the maximality of D_{xy} . On the other hand, from (5.2b) we get $U = Z - D_{xy} \subseteq A_{xy}$ with $q(U) = \nu - 1$.

Assume now $X_0 \subseteq A_{xy}$ is a minimal set with $q(X_0) = \nu - 1$. Since $\nu \geq 2$, there exists an 12-edge $wz \in I_F(X_0)$. Let us choose this with D_{wz} maximal, or equivalently, A_{wz} minimal. By Lemma 5.8, $A_{wz} \subseteq X_0$. Choose a minimal $Y_0 \subseteq Y$ with $q(Y_0) = \nu - 1$, and let $uv \in I_F(Y_0)$ be an arbitrary 12-edge. By Lemma 5.17, w and u are light endnodes of wz and uv , respectively. Consider the hexa-flipping of (xy, uv, wz) . This decreases $q(X)$ and $q(Y)$ by 2; hence by the extreme choice of F , there exists an increasing set Z with $q(Z) \geq \nu - 2$. Let $T = \{x, y, u, v, w, z\} \cap Z$. By possibly complementing Z , we get that T is one of the six sets in Lemma 5.31. Assume Z is chosen minimal.

- (I) T is one of $\{x, v\}$, $\{x, v, w\}$ and $\{x, v, z\}$. If (5.2a) held for X and Z , then $q(X \cup Z) = \nu - 1$. This is a contradiction since u is the light endnode of uv and $V - (X \cup Z) \subseteq Y$ is a $u\bar{v}$ -set. However, (5.2b) gives $q(X - Z) = \nu - 1$, a contradiction to the maximality of D_{xy} , since $X - Z$ is an $\bar{x}y$ -set containing at least one of w and z .
- (II) $T = \{u, z\}$ or $T = \{u, z, x\}$. Since u is the light endnode of uv , (5.2a) cannot hold for Y and Z , thus we may apply (5.2b), yielding $q(Z - Y) = \nu - 1$. This contradicts Lemma 5.8 since $z \in X_0 \cap (Z - Y)$ and $y \in X - (X_0 \cup (Z - Y))$.
- (III) $T = \{w, y\}$. Let us consider the three sets X_0 , Z and $Z' = Z_{vw}$. We use an argument similar to the one in the proof of Lemma 5.19. By the minimal choice of A_{wz} , Claim 5.18 is applicable, and thus $I_F(A_{wz}) = 0$. We shall prove that (5.2a) does not hold for any two of the three sets X_0 , Z and Z' , and the intersection of any two of them is A_{wz} . These easily imply that A_{wz} is a marginal set. The proof is illustrated in Figure 5.2.




Figure 5.2: Illustration of the proof of Lemma 5.32(i) for $T = \{w, y\}$.

First, consider X_0 and Z' . By Lemma 5.16, $Z' \cap X = A_{wz}$; this implies $Z' \cap X_0 = A_{wz}$ and $y \notin Z'$. By the minimality of X_0 , $q(Z' \cap X_0) \leq \nu - 2$; and by Claim 5.5, $q(X_0 \cup Z') \leq \nu - 2$, hence (5.2a) cannot hold for Z' and X_0 . Next, we show that (5.2a) cannot hold for X_0 and Z either. $q(Z \cap X_0) \leq \nu - 2$ again by the minimal choice of X_0 . If $x \in X_0$, then $q(Z \cup X_0) \geq \nu + 1$ since $d_F(Z, X_0) \geq 1$, yielding a contradiction. If $x \notin X_0$, then we get $q(Z \cup X_0) = \nu - 1$ and $Z \cup X_0$ is an $\bar{x}y$ -set, contradicting the maximality of D_{xy} .

Finally, assume (5.2a) were true for Z and Z' . We claim that $q(Z \cap Z') \leq \nu - 2$. This is trivial by Claim 5.5 if $(Z \cap Z') \cap Y \neq \emptyset$. On the other hand, if $Z \cap Z' \subseteq X$, then this follows by Lemma 5.8, since $u \in X_0 \cap (Z \cap Z')$ and $y \in X - (X_0 \cup (Z \cap Z'))$. Consequently, $q(W) = \nu - 2$ for $W = Z \cup Z'$ and $x \notin W$. (This follows since $x \in W$ would imply $d_F(Z, Z') \geq 1$, yielding $q(W) = \nu$.) A similar argument to the one above for X_0 and Z shows that (5.2a) cannot hold for X_0 and W . However, (5.2b) gives $q(W - X_0) \geq \nu - 1$, a contradiction to Claim 5.5 since $v, y \in W - X_0$. It may also be easily verified that A_{wz} is the intersection of any two of the sets X_0 , Z and Z' ; we leave this to the reader.

For the rest, $X \cap Z_{xv} = A_{xy}$ follows by Lemma 5.16. If D_{uv} and Z_{xv} are crossing, then (5.2b) cannot hold for D_{uv} and Z_{xv} , while (5.2a) gives $q(D_{uv} \cup Z_{xv}) = \nu - 2$.

(ii) We start by proving $A_{wz} \subseteq A_{xy}$, or equivalently, $D_{wz} \supseteq D_{xy}$. Assume D_{wz} and D_{xy} are crossing ($D_{wz} \subseteq D_{xy}$ is excluded by $z \in D_{wz} - D_{xy}$). (5.2b) gives a contradiction, since part (i)

implies $q(D_{wz} - D_{xy}) \leq \nu - 2$. (5.2a) is also impossible, since $D_{xy} \cup D_{wz}$ is a \overline{wz} -set, and thus the maximality of D_{wz} implies $q(D_{xy} \cup D_{wz}) \leq \nu - 2$.

As in the proof of part (i), let us choose an 12-edge $uv \in I_F(Y)$; we already know that its light endnode is u . Consider the hexa-flipping of (xy, uv, wz) . We get a set Z with $q(Z) \geq \nu - 2$ and with six possible sets T as in the first part. Case (I) is settled by an identical argument. In the case (II), the existence of the set $Z - Y$ with $q(Z - Y) = \nu - 1$, $y \notin Z - Y$ and $(Z - Y) \cap A_{xy} \neq \emptyset$ gives a contradiction to part (i). Let us now turn to case (III); assume again Z is chosen minimal.

We claim that $Z \subseteq X$. Indeed, if Z and X were crossing and (5.2a) held, then $q(Z \cap X) \leq \nu - 3$ by the minimality of Z , leading to contradiction. If (5.2b) held, then $q(X - Z) \leq \nu - 2$ by part (i) and thus $q(Z - X) = \nu$, a contradiction again.

Consider the set $Z' = Z_{vw}$. By part (i), $Z' \cap X = A_{wz} \subseteq A_{xy}$, thus Z' and Z are crossing. We claim that (5.2b) must hold for Z and Z' . For a contradiction, assume (5.2a) held for them. Then $q(Z \cap Z') \leq \nu - 2$ by part (i), and thus $q(W) = \nu - 2$ for $W = Z \cup Z'$, furthermore, $x \notin Z'$ (as $x \in Z'$ would give $d_F(Z, Z') \geq 1$). (5.2a) for X and W is impossible since it would give $q(V - (X \cup W)) = \nu - 1$, a contradiction as it is an $u\overline{v}$ -subset of Y , and u is the light endnode of uv . On the other hand, (5.2b) implies $q(X - W) = \nu - 1$, a contradiction as x is the light endnode of xy .

For Z and Z' , (5.2b) gives $q(Z - Z') = q(Z' - Z) = \nu - 1$, $\overline{d_{G+F}}(Z, Z') = 1$. Part (i) and the maximal choice of D_{xy} implies that $Z - Z' \subseteq D_{xy}$ and $Z \cap Z' = A_{wz}$. This yields $d_{G+F}(A_{wz}, A_{xy} - A_{wz}) = 1$, as required. Also, (5.2b) cannot hold for D_{xy} and Z ; hence $q(D_{xy} \cup Z) = \nu - 2$. The proof is complete since $D_{xy} \cup A_{wz} = D_{xy} \cup Z$.

(iii) is a trivial consequence of Claim 5.9 for D_{xy} and D_{wz} and the steadiness of A_{xy} and A_{wz} .

(iv) For a contradiction, assume that $w \in D_{xy}$ or $w \in Y$. The first case contradicts part (i): although w is the light endnode of wz , D_{xy} is a $w\overline{z}$ -set with $q(D_{xy}) = \nu - 1$. Hence $w \in Y$; let $uv \in I_F(Y)$ be an arbitrary 12-edge, and consider the hexa-flipping (xy, uv, wz) . There must be an increasing set Z as in the proofs of (i) and (ii), and we examine the same cases (I)-(III). In each case, $q(Z) = \nu - 2$ as $q(Z) = \nu - 1$ is excluded by Claim 5.5; let us choose Z minimal.

(I) Z is a minimal (and stable) $xv\overline{wy}$ -set with $q(Z) = \nu - 2$ and thus $Z = Z_{xv}$. By part (i), $Z_{xv} \cap X = A_{xy}$ and hence $z \in Z$. Consequently, $w \notin Z$. (5.2b) is impossible for Z and Y because x is a light endnode of xy . (5.2a) cannot hold either, since $d_{G+F}(Z, Y) \geq 1$ because of the edge wz .

(II) Since u is the light endnode of uv , we get $q(Z - Y) = \nu - 1$ by (5.2b). This contradicts part (i) since $Z - Y$ intersects A_{xy} and $y \notin Z - Y$.

(III) For X and Z , (5.2a) is impossible since $d_{G+F}(X, Z) \geq 1$. On the other hand, (5.2b) cannot hold either since x is the light endnode of xy .

□

Some prerequisites are needed to prove Lemma 5.34. $W \subseteq X$ is called a **witness set** for X if there exists sets $A_{x_1y_1}, \dots, A_{x_\delta y_\delta} \in \mathcal{A}^0$ (with $\delta \geq 1$) so that $\bigcup_{i=1}^w A_{x_i y_i} = W$, and $(\bigcup_{i=1}^{j-1} A_{x_i y_i}) \cap A_{x_j y_j} \neq \emptyset$ for $j = 2, \dots, \delta$. $A_{x_1y_1}, \dots, A_{x_\delta y_\delta}$ is called a **construction sequence** for W . Note that witness sets are exactly the node sets of connected subhypergraphs of the hypergraph (X, \mathcal{A}_0) . Witness sets for Y are defined analogously.

Lemma 5.36. (i) Every witness set W for X is steady, $q(X - W) = \nu - 1$ and $d_F(W, X - W) > 0$.

(ii) If $wz \in I_F(W)$, then $wz \in I_F(A_{x_i y_i})$ for some member of the construction sequence.

(iii) If W and W' are two witness sets for X , then $d_{G+F}(W, W') = 0$. If A is a witness set for X and B is a witness set for Y , then they satisfy (5.4).

Proof. (i) Consider a construction sequence for W as in the definition. If $\delta = 1$, then we are done by Lemma 5.32(i). Assume now $\delta > 1$. By induction, $W' = \bigcup_{i=1}^{\delta-1} A_{x_i y_i}$ is a steady set with $q(X - W') = \nu - 1$. Let $A = A_{x_\delta y_\delta}$, $D = X - A = D_{x_\delta y_\delta}$. We may assume that D and $X - W'$ are crossing, as otherwise $W = W'$ or $W = A$ or we get a contradiction to the stability of A and W' . The stability also excludes (5.2b) for $X - W'$ and D . (5.2a) implies $q(D \cup (X - W')) = q(D \cap (X - W')) = \nu - 1$, $d_{G+F}(A, W') = 0$. Since $X - W = D \cap (X - W')$, it remains to prove the steadiness of W .

Indeed, assume there is a set $U \subseteq W$ with $q(U) = \nu - 1$. As W' and A are steady, both sets $U \cap (W' - A)$ and $U \cap (A - W')$ are nonempty. (5.2b) cannot hold for U and D , since $q(U - D) \leq \nu - 2$ by the stability of A . (5.2a) implies $U \cup D = X$, $d_{G+F}(U, D) = 0$. These, together with $d_{G+F}(A, W') = 0$ yield $x_\delta \in A \cap W'$, $y_\delta \in W' - A$. By the induction hypothesis, $x_\delta y_\delta \in I_F(A_{x_i y_i})$ for some $i < \delta$. Then, by Lemma 5.32(ii), $A \subseteq A_{x_i y_i} \subseteq W$, a contradiction.

(ii) follows by $d_{G+F}(A, W') = 0$ and the inductual hypothesis. The first part of (iii) is immediate by Lemma 5.32(iii). For the second part, if $A = A_{xy}$ and $B = A_{uv}$ for 12-edges $xy \in I_F(X)$, $uv \in I_F(Y)$, then Lemma 5.32(i) proves (5.4). For larger witness sets, it can be verified easily by induction as in part (i). □

Let us now prove further useful properties of witness sets. The next claim is straightforward by the definition of D_{xy} .

Claim 5.37. If for an 12-edge $xy \in I_F(X)$, W is a witness set and also an $x\bar{y}$ -set, then $A_{xy} \subseteq W$. □

Lemma 5.38. Let U, T and W be three witness sets for X . Assume that T is a descendant of U , and let $wz \in F$ be the unique 12-edge from T to $U - T$.

(i) If $T \cap W \neq \emptyset$ and $z \notin W$, then $U \cap W \subseteq T$. Consequently, $T \cup W$ is a descendant of $U \cup W$.

(ii) If T is also a descendant of W or $T \cap W = \emptyset$, then T is a descendant of $U \cup W$.

(iii) If $W \in \mathcal{A}_0$, then T is always a descendant of $U \cup W$ whenever the condition in part (i) is not met.

Proof. (i) Consider the sets $Z = T \cup (X - U)$ and $D = X - W$. $q(Z) = \nu - 2$ as T is a descendant of U , and $q(D) = \nu - 1$ by Lemma 5.36(i). For Z and D , (5.2b) is impossible because of $q(Z - D), q(D - Z) \leq \nu - 2$, as $Z - D$ and $D - Z$ are nonempty subsets of the steady sets W and U , respectively. (The nonemptiness follows since $T \cap W \subseteq Z - D$ and $z \in D - Z$.) (5.2a) yields $q(D \cap Z) + q(D \cup Z) - 2d_{G+F}(D, Z) \geq 2\nu - 3$. The proof is illustrated in Figure 5.3.




Figure 5.3: Illustration of the proof of Lemma 5.38(i); the sets Z and D are vertically and horizontally striped, respectively.

If $q(D \cup Z) = \nu$, then $D \cup Z = X$, or equivalently, $U \cap W \subseteq T$, as required. This is always the case if $w \in W$ as it implies $d_{G+F}(D, Z) \geq 1$.

Let us assume $q(D \cup Z) \leq \nu - 1$, and thus $w \in T - W$ and $q(D \cap Z) \geq \nu - 2$. Let $Z' = D \cap Z$ and $D' = X - T$. Note that $w \in Z'$. For Z' and D' , (5.2b) is again impossible: $Z' - D'$ and $D' - Z'$ are subsets of the steady sets T and $U \cup W$, respectively. Thus (5.2a) must hold. $d_{G+F}(Z', D') \geq 1$ because of the edge wz . Consequently, $q(Z' \cap D') + q(Z' \cup D') \geq 2\nu - 1$, a contradiction as both are proper subsets of X .

The last part follows since we have just proved $(U \cup W) - (T \cup W) = U - T$. Since $d_{G+F}(U, W) = 0$ by Lemma 5.36(iii), this also implies that wz is the only edge in $G + F$ between $T \cup W$ and $(U \cup W) - (T \cup W)$.

(ii) Consider the sets $T \cup (X - U)$ and $T \cup (X - W)$ if T is also a descendant of W ; and $T \cup (X - U)$ and $X - W$ in case of $T \cap W = \emptyset$. In both cases, (5.2b) is excluded by steadiness, while (5.2a) yields $q(T \cup (X - (U \cup W))) = \nu - 2$. (Notice that $q(T \cup (X - (U \cup W))) \geq \nu - 1$ is impossible since w is the light endnode of wz . $q(T \cup (X - (U \cap W))) \leq \nu - 2$ similarly follows in the first case.) $d_{G+F}(T, (U \cup W) - T) = 1$ is trivial in the first case, whereas it follows by Lemma 5.36(iii) in the second case.

(iii) Let $W = A_{xy}$ for some $xy \in I_F(X)$. Using part (ii), it remains to investigate the case when $T \cap W \neq \emptyset$ and $z \in W$. By Lemma 5.32(iv), $w \in W$, whereas part (ii) of the same lemma implies that A_{wz} is a descendant of W . Furthermore, $A_{wz} \subseteq T$ by Claim 5.37. Let us apply part (i) for W, A_{wz} and T in place of U, T and W , respectively. We get $W \cap T = A_{wz}$ and that $T = T \cup A_{wz}$ is a descendant of $W \cup T$. Now we may apply part (ii) for U, T and $W \cup T$, leading to the desired conclusion. \square

Corollary 5.39. *For any $A_{xy} \in \mathcal{A}^0$, the set system $\mathcal{U} = \{A_{wz} : wz \in I_F(A_{xy})\}$ is laminar. Consequently, if W is a witness set whose construction sequence consists of sets in \mathcal{U} , then $W \in \mathcal{U}$.*

Proof. Indeed, assume $T = A_{wz}$ and $W = A_{w'z'}$ are crossing sets with $wz, w'z' \in I_F(A_{xy})$, both descendants of $U = A_{xy}$. $z \in W$ is impossible as Lemma 5.32(iv) and (ii) would imply $T \subseteq W$, and thus Lemma 5.38(i) is applicable, yielding $W \subseteq T$, a contradiction again. \square

Let us now define an ordering $A_1, A_2, \dots, A_\gamma$ of the elements \mathcal{A}^0 among auxiliary witness sets W_1, \dots, W_γ ($\gamma = |\mathcal{A}^0|$). Let A_1 be an arbitrary minimal element of \mathcal{A}^0 and let $W_1 = A_1$.

In step $i \geq 2$, let $\mathcal{R} = \mathcal{A}^0 - \{A_j : j < i\}$, that is, the sets which have not yet been indexed. Assume first that there exists an $A \in \mathcal{R}$ with $A \cap W_{i-1} \neq \emptyset$. Let us choose such an A minimal for containment, and subject to this, $|A - W_{i-1}|$ minimal. Let $A_i = A$, $W_i = W_{i-1} \cup A_i$. If $A \cap W_{i-1} = \emptyset$ for every $A \in \mathcal{R}$, then let A_i be an arbitrary minimal element of \mathcal{R} and let $W_i = A_i$.

Notice that in this ordering, the connected components of the hypergraph (X, \mathcal{A}^0) will be the maximal W_i 's, and their building sequences are ‘‘continuous’’ subsets of $\{1, \dots, \gamma\}$.

Proof of Lemma 5.34. Let $F_i = \{xy \in I_F(X) : A_{xy} = A_j \text{ for some } j \leq i\}$. Note that $F_\gamma = I_F(X)$. In what follows, we construct a witness system \mathcal{A}_i for F_i consisting of witness sets, providing a witness system \mathcal{A}_γ for $I_F(X)$. A witness system for $I_F(Y)$ can be constructed similarly. Since both consist of witness sets, they are automatically linked by Lemma 5.36(iii), and thus the claim follows.

The members of \mathcal{A}_i will be witness sets whose construction sequences contain only the sets A_1, \dots, A_i . Furthermore, it will be obvious from the construction that \mathcal{A}_i contains all maximal such witness sets (in particular, W_i .) Note that, by the indexing rule, W_{i-1} will be the only

maximal member of \mathcal{A}_{i-1} intersecting A_i . Let $\mathcal{A}_1 = \{A_1\}$. For some $i \geq 2$, assume we have already constructed \mathcal{A}_{i-1} .

(I) If $A_i \cap W_{i-1} = \emptyset$ or $W_{i-1} \subseteq A_i$, then let $\mathcal{A}_i = \mathcal{A}_{i-1} \cup \{A_i\}$. This clearly satisfies the conditions. Note that if $W_{i-1} \subseteq A_i$ then Corollary 5.39 implies that $W_{i-1} \subseteq \{A_1, \dots, A_{i-1}\}$ and thus all these sets are descendants of A_i .

(II) Assume next $A_i \subseteq W_{i-1}$. If \mathcal{A}_{i-1} blocks the entire edge set F_i , then let $\mathcal{A}_i = \mathcal{A}_{i-1}$. Otherwise, there exists an 12-edge $xy \in F_i - F_{i-1}$ not blocked by \mathcal{A}_{i-1} but only by $A_i = A_{xy}$. We shall prove that $\mathcal{A}_i = \mathcal{A}_{i-1} \cup \{A_i\}$ satisfies the conditions.

$xy \in I_F(W_{i-1})$, hence by Lemma 5.36(ii), $xy \in I_F(A_j)$ for some $j < i$. By Lemma 5.32(ii), A_i is a descendant of A_j . The selection rule in step j implies $A_i \cap W_{j-1} = \emptyset$ and $A_j - W_{j-1} \neq \emptyset$. We claim that $W_i = W_j$. Indeed, $A_i \subseteq W_j$, and thus choosing A_ℓ with $A_\ell - W_{\ell-1} \neq \emptyset$ in step $j < \ell < i$ contradicts the selection rule, as $A_i - W_{\ell-1} = \emptyset$. On the other hand, for $j < \ell < i$, either $A_\ell \cap W_{j-1} = \emptyset$ or $A_\ell - W_{j-1} = A_j - W_{j-1}$, as otherwise we would have had a better choice in step j . Together with Corollary 5.39, these guarantee the laminarity of \mathcal{A}_i .

It is left to prove (c) in Definition 5.33. Let C be the smallest member of \mathcal{A}_i containing A_i . Clearly, $C = A_j \cup W$ for some witness set $W \subseteq W_{j-1}$. Lemma 5.38(ii) for $U = A_j$, $T = A_i$ and W gives that A_i is a descendant of C . Next, assume A_i is the smallest set in \mathcal{A}_i containing some $T \in \mathcal{A}_i$. Again, Corollary 5.39 ensures that this is only possible if $T = A_\ell$ for some $\ell < i$, and A_ℓ is a descendant of A_i by Lemma 5.32(ii).

(III) Finally, assume A_i and W_{i-1} are crossing. For an 12-edge $xy \in F$, $A_i = A_{xy}$ implies $y \notin W_{i-1}$ as otherwise Lemma 5.32(iv) and (ii) would give $A_i \subseteq W_{i-1}$. Consequently, xy is also blocked by $W_i = W_{i-1} \cup A_i$. Let $\mathcal{T} \subseteq \mathcal{A}_{i-1}$ denote the set of the largest proper subsets of W_{i-1} . Note that \mathcal{T} forms a subpartition of W_{i-1} , and according to (c) in Definition 5.33, all members of \mathcal{T} are descendants of W_{i-1} . We distinguish three cases. In each of them, we assume that the conditions of the previous case(s) are not met.

(IIIa) There is an 12-edge $wz \in F_{i-1}$ with $w \in A_i \cap W_{i-1}$, $z \in A_i - W_{i-1}$. By Claim 5.37, $A_{wz} \subseteq W_{i-1}$. The conditions in Lemma 5.38(i) are met for A_i, A_{wz} and W_{i-1} in place of U, T and W , thus $W_{i-1} = W_{i-1} \cup A_{wz}$ is a descendant of W_i . Consequently, $\mathcal{A}_i = \mathcal{A}_{i-1} \cup \{W_i\}$ is an appropriate choice.

(IIIb) There is a $B \in \mathcal{T}$ with $B \cap A_i \neq \emptyset$ and $z \notin A_i$, where $wz \in F_{i-1}$ is the unique 12-edge between B and $W_{i-1} - B$. The conditions in Lemma 5.38(i) hold for W_{i-1}, B and A_i , hence $W_{i-1} \cap A_i \subseteq B$ and $A_i \cup B$ is a descendant of W_i . This also implies that all sets in $\mathcal{T} - B$ are disjoint from A_i and hence by Lemma 5.38(ii), they are all descendants of W_i . As the condition in (IIIa) is not met, all edges in F_{i-1} blocked by W_{i-1} are also blocked by W_i , and those blocked by B are also blocked by $A_i \cup B$. Now $\mathcal{A}_i = (\mathcal{A}_{i-1} - \{W_{i-1}, B\}) \cup \{W_i, A_i \cup B\}$ is a witness system for F_i .

(IIIc) Otherwise, Lemma 5.38(iii) yields that all members of \mathcal{T} are descendants of W_i . Setting $\mathcal{A}_i = (\mathcal{A}_{i-1} - \{W_{i-1}\}) \cup \{W_i\}$ satisfies the conditions. \square

So far, we used only Claim 1.18 on the skew supermodularity of R . In the proof of Lemma 5.35 we will however need the stronger Claim 1.19 stating that if (1.5a) or (1.5b) does not hold, then the other holds with equality. Consequently, (5.2a) and (5.2b) have the same property. This will be needed to prove the next claim.

Claim 5.40. *If H_1, \dots, H_δ are disjoint members of \mathcal{A} , then $q(X - \bigcup_{i=1}^\delta H_i) = \nu - \delta$ and $q(\bigcup_{i=1}^\delta H_i) < \nu - \delta$. The same hold for Y and \mathcal{B} .* \square

Proof. We prove the two claims together by induction on δ . For $\delta = 1$ these follow by Lemma 5.36(i); assume we have already proved them for $1, \dots, \delta - 1$. Consider the sets $D = X - \bigcup_{i=1}^{\delta-1} H_i$ and $D' = X - H_\delta$. By induction, $q(D) = \nu - \delta + 1$ and $q(D') = \nu - 1$. (5.2b) cannot hold since $D - D' = \bigcup_{i=1}^{\delta-1} H_i$, $D' - D = H_\delta$, and thus by induction $q(D - D') < \nu - \delta + 1$ and $q(D' - D) < \nu - 1$. Hence (5.2a) holds with equality. Now $q(D \cup D') = \nu$ as $D \cup D' = X$, yielding the first part of the claim.

For the second part, let $Z = \bigcup_{i=1}^\delta H_i$. Assume for a contradiction that $q(Z) \geq \nu - \delta$. Lemma 5.36(i) and (iii) together imply $d_F(Z, D') \geq 1$. If (5.2a) held for Z and D' then we get a contradiction since $q(Z \cup D') = \nu$ and $q(Z \cap D') < \nu - \delta + 1$ by the induction hypothesis. On the other hand, (5.2b) is also impossible since $Z - D' = X - \bigcup_{i=1}^{\delta-1} H_i$ and $D' - Z = H_\delta$. $q(D' - Z) < \nu - 1$ and we have just proved in the first part that $q(Z - D) = \nu - \delta$. \square

Proof of Lemma 5.35. We use Lemma 5.28 to verify that \mathcal{H} is a medusa. For any $1 \leq i, j, t \leq \ell$, $i \neq j$, we construct a separating pair Z, Z' for i and j coherent with t , so that (5.2a) does not hold for them, and $\overline{d}_G(Z, Z') = 0$. Note that this also implies $d_G(C_i, C_j) = 0$ and hence the conditions of the lemma are satisfied. Let A, A' and B be the corresponding members of \mathcal{A} or \mathcal{B} for C_i, C_j and C_t , respectively.

We start by showing the existence of a separating pair (regardless to t). **(I)** First, if $C_i \in \mathcal{C}^X$, $C_j \in \mathcal{C}^Y$, then let $Z = X$, $Z' = A \cup (Y - A')$. $q(Z) = \nu$ and $q(Z') = \nu - 2$ since \mathcal{A} and \mathcal{B} are linked. (5.2a) would contradict the steadiness of A ; in the case of (5.2b), $\overline{d}_{G+F}(Z, Z') = 0$ follows since $q(Z - Z') + q(Z' - Z) \leq 2\nu - 2$.

(II) Let us now assume that C_i and C_j are both in \mathcal{C}^X or both in \mathcal{C}^Y ; wlog. consider \mathcal{C}^X . **(IIa)** If A and A' are disjoint, then let $Z = Y \cup A$, $Z' = X - A'$. $q(Z) = q(Z') = \nu - 1$ (notice that $Z = V - (X - A)$), and the same argument works as in the first case. **(IIb)** If $A \subseteq A'$, then we may assume that A is a descendant of A' (otherwise, we replace A by the largest set A'' with $A \subsetneq A'' \subsetneq A'$). For $Z = A \cup (X - A')$ and $Z' = Y \cup A$ we have $q(Z) = \nu - 2$ and $q(Z') = \nu - 1$. (5.2a) is impossible since $q(Z \cap Z') \leq \nu - 2$ because of $Z \cap Z' = A$, and $q(Z \cup Z') \geq \nu - 2$ as $V - (Z \cup Z')$ is a subset of the steady set A' . From (5.2b) we get $\overline{d}_{G+F}(Z, Z') \leq 1$. However, we know that there exists an 12-edge $wz \in I_F(A')$ from A to $A' - A$, hence $\overline{d}_G(Z, Z') = 0$. **(IIc)** The argument is the same for the case $A' \subseteq A$ by changing the role of A and A' and complementing the sets Z and Z' . (Hence we set $Z = (A - A') \cup Y$ and $Z' = X - A'$.)

We need a separating pair with the stronger property of being consistent with t . Let us reconsider the cases above. **(I)** In the construction above, $A \subseteq Z \cap Z'$ and $A' \cap (Z \cup Z') = \emptyset$. Hence if $B \subseteq A$, then any pair separating i and j is automatically consistent with t . Also, if $A \subseteq B$, then a pair separating t and j also separates i and j , and is consistent with t . By similar arguments, we are also done if $B \subseteq A'$ or $A' \subseteq B$. The remaining case is when B is disjoint from both A and A' . If $B \subseteq Z - Z' = X - A$, then let $\hat{Z} = X - B$. If $B \subseteq Z' - Z = Y - A'$, then let $\hat{Z} = X \cup B$. In both cases, $q(\hat{Z}) = \nu - 1$ and it can be verified easily that \hat{Z}, Z' is an appropriate choice.

(IIa) Again, the nontrivial cases is when B is disjoint from both A and A' . If $B \subseteq Y$ then let $\hat{Z} = A \cup (Y - B)$, $\hat{Z}' = Z'$ and if $B \subseteq X$, then let $\hat{Z} = Z$ and $\hat{Z}' = X - (A \cup B)$. It is easy to show that \hat{Z}, \hat{Z}' is a good pair in both cases, however, Claim 5.40 is needed for the proof.

In the case **(IIb)**, we have to investigate $B \subseteq Y$ and $B \subseteq X - A'$. Let $\hat{Z}' = A \cup (Y - B)$ in the first while $\hat{Z}' = A \cup B \cup Y$ in the second case and $\hat{Z} = Z$ in both cases. It is left to the reader to verify, using Claim 5.40, that \hat{Z}, \hat{Z}' is a good pair. **(IIc)** can be again handled similarly.

Having proved that \mathcal{H} is a medusa, we shall verify $\tau_1(G, r, \vec{m}, \mathcal{H}) = \nu$. Let \mathcal{A}_M and \mathcal{B}_M denote the set of the maximal components of \mathcal{A} and \mathcal{B} , respectively; let $|\mathcal{A}_M| = s$ and $|\mathcal{B}_M| = t$. Furthermore, let $F_1 \subseteq F$ be the set of ordinary edges w.r.t. \mathcal{H} and $h = 1$ (as defined before Lemma 5.25). Let $F_2 = F - F_1$. Let $G' = G + F_1$, and let \vec{m}' denote the ‘‘degree vector’’ of F_2 , that is, (1.9) holds for \vec{m}' and (F_2, φ) .

Notice that $I_F(X^*) = I_F(Y^*) = I_F(C_i) = \emptyset$ for each $C_i \in \mathcal{C}$, and there are exists no 12-edge $xy \in F$ with $x \in X^* \cup Y^*$, $y \in \bigcup \mathcal{C} = T_X \cup T_Y$. The edges in F_2 are exactly those in F connecting two tentacles in \mathcal{C}^X or two in \mathcal{C}^Y . Therefore, $m(V) - m_1(\bigcup \mathcal{C}) = m(X^*) + m(Y^*) + m_2(\bigcup \mathcal{C}) = d_F(X^*) + d_F(Y^*) + |F_2|$. Hence we may rewrite $\tau_1(G, r, \vec{m}, \mathcal{H})$ in the form

$$\tau_1(G, r, \vec{m}, \mathcal{H}) = \frac{1}{2} (\chi_1 + q(X^*) + q(Y^*) - |F_2|).$$

The proof finishes by the following claim.

Claim 5.41. (i) $q(X^*) = \nu - s$, $q(Y^*) = \nu - t$.

(ii) $|\mathcal{C}| = s + t + |F_2|$.

(iii) Every tentacle C_i is 1-odd.

Proof. **(i)** is immediate by Claim 5.40. **(ii)** By Lemma 5.36(iii), $d_{F_2}(T_X, T_Y) = 0$. We claim that $|\mathcal{C}^X| = |\mathcal{A}| = s + |I_{F_2}(T_X)|$ and analogously for \mathcal{C}^Y . Indeed, by (e) in Definition 5.33, there is a unique 12-edge in F_2 between A' and $A - A'$ for each $A' \in \mathcal{A} - \mathcal{A}_M$ with A being the smallest set containing it. Hence there is a bijection between $I_{F_2}(T_X)$ and $\mathcal{A} - \mathcal{A}_M$.

(iii) By Lemma 5.25, the set of 1-odd tentacles is the same for \mathcal{H}, G, \vec{m} and $\mathcal{H}', G', \vec{m}'$, where $\mathcal{H}' = (X^*, Y, \mathcal{C}^X)$. Notice also that $p' = q_{F_1}$, where p' denotes the demand function for G' . For

a tentacle $C \in \mathcal{C}^X$, let A denote the corresponding member in \mathcal{A} . Let $\mathcal{T} \subseteq \mathcal{A}$ denote the set of largest sets contained in A ; let $|\mathcal{T}| = a$. For each $A' \in \mathcal{T}$, $d_{G+F}(A', A - A') = 1$, and the unique edge is an 12-edge $xy \in F_2$ with $x \in A'$, $y \in C$.

$q_{F_1}(Y) = q_F(Y) = \nu$ and $q_F(Y \cup A) = \nu - 1$. Let $b_1 = d_{F_2}(C, X - A)$ and $b_2 = d_{F_2}(\bigcup \mathcal{T}, X - A)$. Note that if $A \in \mathcal{A}_M$, then $b_1 = b_2 = 0$, and if $A \notin \mathcal{A}_M$ then $b_1 + b_2 \geq 1$. Thus $q_{F_1}(Y \cup A) = \nu - 1 + b_1 + b_2$. By Claim 5.40, $q_F(Y \cup (\bigcup \mathcal{T})) = \nu - a$, and thus $q_{F_1}(Y \cup (\bigcup \mathcal{T})) = \nu + b_2$. By the hydra property, $p(Y \cup A) + p(Y) = p(Y \cup (\bigcup \mathcal{T})) + p(Y \cup C)$. Since $d_{F_1}(\bigcup \mathcal{T}, C) = 0$, it follows that

$$q_{F_1}(Y \cup C) - q_{F_1}(Y) = q_{F_1}(Y \cup A) - q_{F_1}(Y \cup (\bigcup \mathcal{T})) = b_1 - 1.$$

$m'_1(C) = b_1$ and thus C is 1-odd for $\mathcal{H}', G', \vec{m}'$ (recall $p' = q_{F_1}$), and consequently, for \mathcal{H}, G, \vec{m} . □

□

5.3.2 The augmentation problem

In this section, we briefly sketch how Conjecture 1.46 could be derived from Conjecture 1.45. We start by constructing a legal degree-prescription $\vec{m} = (m_1, m_2)$ compatible with the partition $\mathcal{Q} = (Q_1, Q_2)$ as in Section 5.2.2. This satisfies $m_1(V) = m_2(V) = \Psi_{\mathcal{Q}}(G)$. Next, consider a \preceq -minimal \vec{m} -prescribed legal edge set F . We are done if $\nu_F = 0$. If $\nu_F > 0$, then consider an optimal hydra $\mathcal{H} = (X^*, Y^*, C_1, \dots, C_\ell)$ and $h \in \{1, 2\}$ with $\nu_F = \tau_1(G, r, \vec{m}, \mathcal{H})$ as in Conjecture 1.46. Wlog. assume $h = 1$. Let $v \in C_i$ with $m_1(v) > 0$; consider the minimum tight set X_v containing v as in Section 5.2.2.

If $X_v \subseteq C_i \cap Q_1$ holds in all such cases, then we may uncross these sets as in Section 5.1.1. This results in a 1-subpartition \mathcal{Z} , which is a refinement of C_1, \dots, C_ℓ and $p(\mathcal{Z}) = m_1(\bigcup C_i)$. Now $\sum(p(Z) : Z \in \mathcal{Z}, Z \subseteq C_i) = m_1(C_i)$ for each $1 \leq i \leq \ell$. Consequently, the 1-odd tentacles are the same as the 1-toxic tentacles, and hence $\tau'_1(G, r, \mathcal{Z}, \mathcal{H}) = \tau_1(G, r, \vec{m}, \mathcal{H}) + \frac{1}{2}m(V) = \nu_F + \frac{1}{2}m(V)$. Finally, Lemma 5.21 yields an augmenting edge set of size $\tau'_1(G, r, \mathcal{Z}, \mathcal{H})$.

If $X_v - (C_i \cap Q_1) \neq \emptyset$ for some $v \in C_i$, $m_1(v) > 0$, then we may define another legal degree-prescription m' and a \preceq -minimal \vec{m}' -prescribed legal edge set F' with $\nu_{F'} < \nu_F$. We do not elaborate this argument here: it needs structural properties of \preceq -minimal legal edge sets generalizing Lemma 5.32. However, these results were proved only under the assumptions $\nu_F \geq 2$ and $\bigcup \mathcal{F}_F = V$.

5.4 Further remarks

Partition-constrained global edge-connectivity augmentation

Let us briefly sketch how the ideas in Section 5.2.1 can be extended to give a new and simpler proof of Theorem 1.23. More precisely, we work here with the degree-prescribed version, which we did not formulate in the thesis. Nevertheless, assume we have a legal degree-prescription \vec{m} so that (1.4) holds, and let us have a connectivity requirement $r \equiv k$. Let F be a \preceq -minimal \vec{m} -prescribed edge set. We shall prove $\nu \leq 1$.

In the case of global connectivity requirements, both (5.1a) and (5.1b) hold for any crossing X, Y with $p(X), p(Y) > 0$. Proving $\nu = 1$ is utterly simple. Indeed, assume $\nu \geq 2$. Consider X_0 as in Lemma 5.8, and $xy \in X_0, uv \in Y$ with (xy, uv) flippable. For X_0 and the stable set Z_{xv} , (5.1b) yields a contradiction. If $\nu = 1$, we can exhibit a C_4 - or C_6 -obstacle⁶ by analyzing a single hexa-flipping.

A similar argument, combined with the ideas of the proof of Theorem 5.11 in Section 5.1.4, could be used to develop a simpler proof of the recent theorem of Bernáth, Grappe and Szigeti [11] on partition-constrained coverings of positively crossing symmetric supermodular functions.

Beyond Theorem 5.30

On the way from Theorem 5.30 towards Conjecture 1.45, the first step would be to leave the assumption $X \cup Y = V$. Lemma 5.32 does not really use this assumption, and remains true with minor modifications. The difficulty comes from the edges incident to $V - (X \cup Y)$. One might give a categorization of such edges, but there is essentially five different types of them. Each type can be characterized in a manner similar to Lemmas 5.6 and 5.32. However, the argument reaches an extreme level of complexity, far beyond the patience of both the author and any possible reader.

To handle edges incident to $V - (X \cup Y)$, we also need a refinement of the partial order \preceq as follows: $F' \prec F$ if $\nu_{F'} < \nu_F$, or $\nu_{F'} = \nu_F$ and $|\mathcal{F}_{F'}| < |\mathcal{F}_F|$, or $\nu_{F'} = \nu_F$ and $|\mathcal{F}_{F'}| = |\mathcal{F}_F|$, but $\sum_{Z \in \mathcal{F}_{F'}} |Z| > \sum_{Z \in \mathcal{F}_F} |Z|$. That is, we also want to maximize $|X| + |Y|$.

For $\nu_F = 1$, the situation is even worse. We needed completely different kind of arguments for $\nu_F = 1$ and $\nu_F \geq 2$ already in the proof of Theorem 1.42. For Conjecture 1.45, we would apparently also need a new type of argument for this case, doubling both length and complexity.

Once having proved Conjecture 1.45, it can be probably easily extended to an arbitrary number of partition classes. For the global connectivity version Theorem 1.23, the main difficulties already occur for $t = 2$. We also need some general version of the C_6 -configuration, but

⁶ C_4 - and C_6 -configurations are for the augmentation problem, while the obstacles for the degree-prescribed problem. Analogously, notice that we also use hydrae in two different senses, with toxic tentacles for the augmentation and odd ones in the degree-prescribed version.

hopefully this is the only new kind of obstacle.

Minimum cost edge-connectivity augmentation problems

Although the minimum-cost version of local edge-connectivity augmentation is NP-complete, however, unlike the other basic connectivity augmentation problems, it admits a nice and strong approximation. Jain [48] proved that for the natural LP-relaxation of the problem, a basic feasible solution always has a component of value at least $\frac{1}{2}$. Rounding up such a value to 1, adding this edge to the graph and iterating the method gives a 2-approximation algorithm.

A natural question is: *for which classes cost functions is local edge-connectivity augmentation polynomially solvable?* An example is - similarly to Chapters 2 and 3 - the class of node-induced cost functions, as it can be shown via standard polyhedral methods. The partition constrained problem can also be interpreted in this framework: given the partition \mathcal{Q} , let $c(uv) = 1$ if u and v lie in different classes of \mathcal{Q} and let $c(uv) = 2$ if u and v are contained in the same class. It is clear that finding a minimum size \mathcal{Q} -legal augmenting edge set is equivalent to finding a minimum cost augmentation, hence the problem for this cost is in P for the global connectivity case - and we conjecture that also for arbitrary requirements.⁷

One might wonder if there is a solvable class containing both node-induced cost functions and the partition-induced cost functions as above. For example, a natural candidate is if we have a different value w_i for each partition class Q_i , and the cost of edges between classes Q_i and Q_j is $w_i + w_j$, while the cost of edges inside class Q_i is $2w_i + 2 \min_{j \neq i} w_j$. (Or equivalently, we want to find a minimum cost \mathcal{Q} -legal augmenting edge set with cost $w_i + w_j$ between Q_i and Q_j . Notice that for this cost function, the cost remains unchanged by a \mathcal{Q} -legal flipping.) We think that this should not be much more difficult than the minimum cardinality partition-constrained problem.

Let us propose another, related question. Jain's iterative rounding method is the only known 2-approximation algorithm for the general minimum cost problem; combinatorial algorithms (e.g. Williamson et. al. [78]) have much worse approximation ratios. A possible approach for constructing a combinatorial 2-approximation could be the following (at least for the uncapacitated case). Find an sufficiently broad class of cost functions \mathcal{K} for which (i) the minimum cost version is still solvable; (ii) arbitrary metric cost function can be 2-approximated by a cost function in \mathcal{K} (that is, for a cost function c , we can find a $c' \in \mathcal{K}$ with $c' \leq c \leq 2c'$). \mathcal{K} being the node-induced cost functions does not meet this latter requirement; however, there might exist a broader class that works. (Nevertheless, partition-induced cost functions should

⁷In these problems, we allow an arbitrary number of copies of the same edge in the augmenting set. In this case, it may always be assumed that the cost function satisfies the triangle inequality. If capacities are also imposed, the problem becomes NP-complete even in the minimum cardinality case (that is, if $c \equiv 1$), as shown by Jordán [51]. Nevertheless, the approximation result of Jain also works with capacities.

be rather excluded from \mathcal{K} : it would be desirable to find a class where a relatively simple algorithm yields an optimal solution.)

Chapter 6

Constructive characterization of (k, ℓ) -edge-connected digraphs

This chapter is devoted to the proof of Theorem 1.47, based on our joint paper [56] with Erika Renáta Kovács. In Section 6.1, the precise definitions are given and some basic properties are exhibited. We also give the proof of Theorem 1.47 here based on the main technical tool Theorem 6.1. This is a special case of the stronger Theorem 6.7 that we prove in Section 6.2 by using three basic lemmas. Among these, the first is a general splitting off result proved in Section 6.3, while the proof of the other two lemmas is given in Section 6.4. Finally, in Section 6.5 we describe the structure of locally admissible sets and present a polynomial algorithm for finding a sufficient locally admissible set F at a special node z . We also show an example of an insufficient maximal globally admissible edge set.

6.1 Basic concepts and the proof of Theorem 1.47

We start with recalling some definitions from Section 1.5.4. Let $D = (V, A)$ be a (k, ℓ) -edge-connected directed graph with root $r_0 \in V$. For $X \subseteq V$, let $\gamma(X) = k$ if $r_0 \notin X$ and $\gamma(X) = \ell$ if $r_0 \in X$. A node $v \in V$ is called **special** if $\rho(v) = k$, $\ell \leq \delta(v) \leq k - 1$. Let S denote the set of special nodes ($S \neq \emptyset$ is not assumed). If $X \subseteq S$ then we say that X is a **special set**. Observe that $r_0 \notin S$ as $\delta(r_0) \geq k$. For a $z \in S$, a subset F of edges entering z is **locally admissible at z** if $D - F$ is (k, ℓ) -edge-connected in $V - z$ and $|F| \leq k - \delta(z)$. A locally admissible F will be called **sufficient** if $|F| = k - \delta(z)$. Theorem 1.47 will be an easy consequence of the following.

Theorem 6.1. *In a minimally (k, ℓ) -edge-connected digraph $D = (V, A)$ there exists a special node z with a sufficient locally admissible set at z .*

Let us see how Theorem 1.47 follows from this.

Proof of Theorem 1.47. First let us show that the operations (i) and (ii) preserve (k, ℓ) -edge-connectivity. This is straightforward in the case of (i). For (ii), let $D' = (V + z, A')$ denote the digraph resulting from the (k, ℓ) -edge-connected digraph $D = (V, A)$ by applying (ii). For every $v \in V - r_0$, the k edge-disjoint paths from r_0 to v and the ℓ edge-disjoint paths from v to r_0 in D naturally give the same number of paths in D' . Thus the only problem could be if there were too few paths from r_0 to z or from z to r_0 .

In this case, by Menger's theorem we have a subset X of $V + z$ with $r_0 \notin X$, $z \in X$, and either $\rho(X) < k$ or $\delta(X) < \ell$. Since D' is (k, ℓ) -edge-connected in V , the only possibility is $X = \{z\}$. However, $\rho(z) = k$ and $\delta(z) \geq \ell$ gives a contradiction.

For the other direction, if D is not minimally (k, ℓ) -edge-connected, then we can obtain D from a smaller (k, ℓ) -edge-connected graph by operation (i). Otherwise, Theorem 6.1 is applicable. Consider the special node z and the sufficient locally admissible F . $D - F$ is (k, ℓ) -edge-connected in $V - z$ and $\rho(z) = \delta(z)$, satisfying the conditions of Theorem 1.34. For the digraph D' resulting by a complete splitting at z , operation (ii) can be applied to get D . \square

The locally admissible edge sets are characterized by the following claim. Let $\Delta^{in}(Z)$ and $\Delta^{out}(Z)$ denote the sets of edges entering and leaving the set Z , respectively. As before, z sometimes stands for the set $\{Z\}$.

Claim 6.2. $F \subseteq \Delta^{in}(z)$ is locally admissible at z if and only if $|F| \leq k - \delta(z)$ and for each $\emptyset \neq X \subsetneq V$, $X \neq \{z\}$,

$$\rho_{A-F}(X) \geq \gamma(X). \tag{6.1}$$

Proof. If F is locally admissible then for $X \neq V - z$, (6.1) is the necessary cut condition as $D - F$ is (k, ℓ) -edge-connected in $V - z$. If $X = V - z$ then it is equivalent to $\delta_{A-F}(z) \geq \ell$, which follows since $\delta_F(z) = 0$. The converse direction follows by Menger's theorem. \square

It is easy to check in polynomial time whether a set of edges entering z is locally admissible. Furthermore these edge sets admit a nice structure: they form a matroid. A consequence is that a building sequence can be found in polynomial time for a (k, ℓ) -edge-connected digraph D . This will be discussed in Section 6.5.

Given an arbitrary edge set $F \subseteq A$, for a node $v \in V$ we use the notation $F_v = F \cap \Delta^{in}(v)$. Let $\mu(X) = \delta_F(V - S - X, X)$, and let $t(X) = \min\{\delta_F(V - S - X, v) : v \in X\}$. A v giving the minimum value in the definition of $t(X)$ is called a **seed** of X . Let $T(X) = \max\{\rho_{F_v}(X) : v \in X\}$, and a v giving the maximum value is called a **sprout** of X . Note that a set may have multiple seeds and sprouts.

Definition 6.3. In a digraph $D = (V, A)$ with special nodes $S \subseteq V$, we say that $F \subseteq A$ is **globally admissible** if

$$\rho(X) \geq \gamma(X) + \rho_F(X), \quad \text{if } X - S \neq \emptyset, X \subsetneq V, \quad (6.2a)$$

$$\rho(X) \geq k + T(X), \quad \text{if } X \text{ is special, } |X| \geq 2, \quad (6.2b)$$

$$\rho(X) \geq \gamma(X) + \mu(X) - t(X), \quad \text{for every } \emptyset \neq X \subsetneq V, \quad (6.2c)$$

$$|F_v| \leq k - \delta(v), \quad \text{for every special node } v \text{ and,} \quad (6.2d)$$

$$F_v = \emptyset, \quad \text{if } v \notin S. \quad (6.2e)$$

Note that if X is not special, then all nodes in $X - S$ are seeds and $t(X) = 0$, and thus (6.2a) implies (6.2c). For a special set X , we have two conditions. On the right hand side of (6.2c), we consider only edges coming from non-special nodes, however, not all such edges are taken into account. The importance of (6.2b) is revealed by the following claim.

Claim 6.4. *If F is globally admissible, then for each $v \in S$, F_v is locally admissible at v .*

Proof. We have to verify (6.1). If X is not special, then $\rho_{A-F_v}(X) \geq \rho_{A-F}(X) \geq \gamma(X)$ by (6.2a). If X is special and $|X| \geq 2$, then by (6.2b), $\rho_{A-F_v}(X) \geq \rho(X) - T(X) \geq k$. \square

Claim 6.5. *If F is globally admissible in D and $F' \subseteq F$, then F' is also globally admissible in D .*

Proof. When removing an edge from F , the right hand sides of (6.2a), (6.2b) and (6.2c) cannot increase. \square

$F = \emptyset$ is globally admissible if and only if D is (k, ℓ) -edge-connected. By the above claim, any digraph D that admits a globally admissible F is automatically (k, ℓ) -edge-connected.

We say that a globally admissible set F is **maximal** if there is no edge $uv \in A - F$ so that $F + uv$ is also globally admissible. A globally admissible F is called **sufficient** if (6.2d) holds with equality for at least one special v , otherwise it is **insufficient**.

Let us now introduce now the various types of tight sets. We say that a set X is **tight** with respect to the globally admissible F if at least one of (6.2a), (6.2b) or (6.2c) holds with equality for X . A tight set with $X - S \neq \emptyset$ is called **normal tight**. A special tight X with $|X| \geq 2$ is called **T -tight** or **μ -tight** if it satisfies (6.2b) or (6.2c) with equality, respectively. For a tight X , if $r_0 \notin X$, then X is called **in-tight**, and if $r_0 \in X$, then $V - X$ is called **out-tight**. Note that, somewhat confusingly, an out-tight set is not necessarily tight.

Claim 6.6. *If F is insufficient globally admissible and for $uv \in A - F$, $v \in S$, $F + uv$ is not globally admissible, then uv enters a tight set X satisfying one of the following: (a) X is a normal tight set, or (b) X is a T -tight set with sprout v , or (c) X is μ -tight, $u \in V - S$ and X has a seed t with $t \neq v$.*

Proof. By assumption, $F + uv$ should violate one of (6.2a), (6.2b) or (6.2c). This cannot happen if none of them holds with equality for F , since the right hand sides may increase by at most 1. Thus uv must enter a tight set X . If X is T -tight and v is not a sprout of v , then $T(X)$ does not increase by adding uv to F and thus (6.2b) will not be violated for X . Similarly, if X is μ -tight and $u \in S$, then (6.2c) remains unchanged for $F + uv$. If $u \notin S$ but the unique seed of X is v , then for $F + uv$, both $\mu(X)$ and $t(X)$ increase by 1. \square

Note that if F is insufficient maximal globally admissible, this claim applies for every edge $uv \in A - F$, $v \in S$.

We will prove a slight generalization of Theorem 6.1 for the purpose of a special induction argument. To formulate this, one more new notion is needed. A globally admissible edge set F **saturates** the digraph D if every edge $uv \in A - F$ with $v \notin S$ enters a normal tight set. We are going to prove the following:

Theorem 6.7. *Let $F_0 \subseteq \Delta^{out}(r_0)$ be an arbitrary globally admissible set of edges in $D = (V, A)$ so that F_0 saturates D . Then there exists a sufficient globally admissible F with $F \supseteq F_0$.*

The (k, ℓ) -edge-connectivity of D is tacitly implied by the existence of F_0 . However, D is not assumed to be minimal subject to this property. Nevertheless, $F_0 = \emptyset$ is a globally admissible edge set saturating D if and only if D is a minimally (k, ℓ) -edge-connected digraph, and thus Theorem 6.1 is a direct consequence of Theorem 6.7. Unfortunately, it is not true that every maximal globally admissible F with $F \supseteq F_0$ is sufficient, as shown by a counterexample in Section 6.5.

Let uv be an edge entering the tight set X . If $v \in S$ and X and uv satisfy one of the conditions in Claim 6.6 or $v \notin S$ and X is normal tight, then we say that X **blocks** uv .

We conclude this section with some elementary propositions.

Claim 6.8. *If $X, Y \subseteq V$, then*

$$\rho(X) + \rho(Y) = \rho(X \cap Y) + \rho(X \cup Y) + d(X, Y), \text{ and} \quad (6.3a)$$

$$\rho(X) + \rho(Y) = \rho(X - Y) + \rho(Y - X) + \rho(X \cap Y) - \delta(X \cap Y) + \bar{d}(X, Y). \quad (6.3b)$$

\square

Claim 6.9. *For any $X, Y \subseteq V$,*

$$\gamma(X) + \gamma(Y) = \gamma(X \cup Y) + \gamma(X \cap Y), \text{ and} \quad (6.4a)$$

$$\gamma(X) + \gamma(Y) \leq \gamma(X - Y) + \gamma(Y - X). \quad (6.4b)$$

\square

Claim 6.10. *For any $X \subseteq V$, $\rho(X) - \delta(X) = \sum_{v \in X} (\rho(v) - \delta(v))$.* \square

Claim 6.11. *Assume F is insufficient globally admissible, and $Z \neq \emptyset$ is special. Then $\delta(Z) < \rho_{A-F}(Z)$.*

Proof. For each $v \in Z$, $\rho(v) - \delta(v) > |F_v|$, and thus by summing for all $v \in Z$, $\rho(Z) - \delta(Z) = \sum_{v \in Z} (\rho(v) - \delta(v)) > \sum_{v \in Z} |F_v| \geq \rho_F(Z)$, hence the claim follows. \square

Claim 6.12. *For $D = (U + u, A)$ with $\rho(u) = \delta(u)$, let D_u denote the result of an (arbitrary) complete splitting at u . Then for any $X \subsetneq U + u$, $\rho_{D_u}(X - u) \leq \rho_D(X)$.*

Proof. If $u \notin X$, then the claim follows since splitting off a pair of edges incident to u cannot increase the degree of $X = X - u$. In the case of $u \in X$, $\rho_{D_u}(X - u) \leq \delta_D(U - X, u) + \delta_D(U - X, X - u) = \rho_D(X)$. \square

6.2 Proof of Theorem 6.7

The proof relies on three basic lemmas. First:

Lemma 6.13. *Let $F_0 \subseteq \Delta^{out}(r_0)$ be an insufficient globally admissible set of edges, and $\rho(u) = \delta(u)$ for some $r_0 \neq u \in V$. There exists a complete splitting at u so that F_0 is globally admissible in the resulting digraph.*

Lemma 6.14. *Assume F' is a globally admissible edge set and X is a tight set with $|X| \geq 2$, $r_0 \notin X$, $|X - S| \leq 1$. Then for any maximal globally admissible $F \supseteq F'$, F is sufficient.*

Lemma 6.15. *If F is maximal globally admissible with $u \in S + r_0$ for each $uv \in F$, then F is sufficient.*

The first of these will be proved in Section 6.3, while the last two in Section 6.4. Let us now turn to the proof of Theorem 6.7. Consider a counterexample $D = (V, A)$ and F_0 so that $|V|$ is minimal, and subject to this, $|F_0|$ is maximal. Consider a maximal globally admissible $F \supseteq F_0$. By the assumption, F is insufficient.

Case I

Assume there is a $u \in V$ with $\rho(u) = \delta(u) = k$. By Lemma 6.13, there is a complete splitting at u so that F_0 is globally admissible in the resulting digraph $D_u = (V - u, A')$.

Claim 6.16. *F_0 saturates D_u .*

Proof. The set of special nodes is the same S in D and D_u . Consider an edge $e = yz$ in D_u with $z \notin S$. Assume first that e is an edge in D as well. There is a normal tight set $X \subseteq V$ blocking e in D , since F_0 saturated D . Claim 6.12 implies $\rho_{D_u}(X - u) \leq \rho_D(X)$. $X - u$ is also

normal and as the subset of F_0 entering $X - u$ in D_u is the same as the subset in D entering X , it follows that $X - u$ blocks e in D_u .

If $e = yz$ is a new edge, then take a set X that blocked uz in D . X is again a normal tight set in D_u . Note that $y \notin X$ as otherwise the in-degree of X would be smaller in D_u than in D while the value of $\rho_{F_0}(X)$ does not change. Hence X blocks e in D_u , completing the proof. \square

As D_u has less nodes than D , by the minimality of $|V|$ there exists a special node w and a sufficient locally admissible edge set F_w so that $F' = F_w \cup F_0$ is globally admissible. Note that w is special in D as well.

From D_u we can get to D by pinching the k splitted edges with u . By abuse of notation, we will denote by F_w the edge set in D corresponding to F_w in D_u in the sense that if an edge $xw \in F_w$ has been divided by u , then we replace xw by uw in F_w . We will also use F' in this sense in D . Unfortunately, it might happen that F' is not globally admissible in D . Consider a globally admissible F_1 maximal subject to the condition $F_0 \subseteq F_1 \subseteq F'$ with $|F_1|$ as large as possible. If $F_1 = F'$, then F_1 is sufficient as $\delta_D(w) = \delta_{D_u}(w)$. Otherwise, we are going to prove that there is a tight set Z for F_1 with $|Z - S| \leq 1$, $|Z| \geq 2$ so Lemma 6.14 is applicable giving a sufficient globally admissible superset of F_1 .

Assume $F_w - F_1 \neq \emptyset$, and consider an edge $zw \in F_w - F_1$. By Claim 6.6, zw is blocked by some tight set Z with respect to F_1 .

Claim 6.17. $Z \subseteq S \cup \{u\}$

Proof. $Z = V - u$ is impossible as $\delta_{F_1}(u) < |F_w| \leq k - \ell$, and thus $\rho_{A-F_1}(V - u) > \ell$. Assume $V - Z - u \neq \emptyset$ and $Z - S - u \neq \emptyset$. As F' is admissible in D_u and $Z - u$ is not special, $\rho_{D_u, A'-F'}(Z - u) \geq \gamma(Z)$ follows. Claim 6.12 implies $\rho_{D, A-F'}(Z) \geq \rho_{D_u, A'-F'}(Z - u)$. However, $\rho_{A-F_1}(Z) > \rho_{A-F'}(Z) \geq \gamma(Z)$ as $zw \in F_1 - F'$ enters Z , showing that Z cannot be tight in D . This implies the claim. \square

Case II

Assume the condition of Case I does not hold and there is an edge $uv \in F$ with $u \in V - S - r_0$. Let $D_1 = (V, A - uv + r_0v)$ and $F_1 = F_0 + r_0v$.

Claim 6.18. F_1 is globally admissible in D_1 and saturates it. The set of tight sets is the same in D and in D_1 .

Proof. If $v \notin X$ or $v \in X$ and $|\{u, r_0\} \cap X| \neq 1$ then no term is changed in the conditions (6.2a), (6.2b) and (6.2c). This is in fact always the case for (6.2b). If $u, v \in X$, $r_0 \notin X$, then in (6.2a) and (6.2c), both sides increase by one, while if $v \in X$, $u \notin X$, $r_0 \in X$, both sides decrease by one. (Note that $t(X) = 0$ in both cases as $X - S \neq \emptyset$.)

This implies the admissibility and that the set of tight sets coincide in the two cases. Thus if an edge $uv \in A - F$ with $v \notin S$ is blocked by a normal tight set for F_0 in D , then the same set blocks it in D_1 , proving the saturation. \square

By the choice of D and F_0 , there is a sufficient edge set $F' \supseteq F_1$ in D_1 with $|F'_w| = k - \delta_{D_1}(w)$ for some w special node in D_1 . All nodes but u and r_0 have the same in- and out-degrees in D and D_1 , and thus w is special in D unless $w = u$ and $\rho(w) = \delta(w) = k$. This is a contradiction since we assumed that no such node exists.

Let $F'' = F' - r_0v + uv$. By the previous claim, it is straightforward to show that F'' is globally admissible in D containing F_0 .

Case III.

For all edges in $uv \in F$, $u \in S + r_0$. The conditions of Lemma 6.15 are satisfied, showing that F is sufficient.

6.3 Splitting off

Theorem 1.1 gave the minimum number of edges covering a positively crossing supermodular function on set pairs. What we are now interested in is an easier problem, namely, coverings of positively crossing supermodular set functions. The following theorem of Frank can be seen as a corollary of Theorem 1.1 on the one hand, and as an abstract generalization of Mader's splitting off theorem (Theorem 1.28) on the other hand.

Analogously as in Section 1.3, we introduce the notion of degree prescribed edge sets in directed graphs. For a ground set U , let us call the pair (m_i, m_o) a **degree prescription** if m_i and m_o are two $U \rightarrow \mathbb{Z}_+$ functions with $m_i(U) = m_o(U)$. We say that H is an (m_i, m_o) -**prescribed edge set** if $\rho_H(v) = m_i(v)$, $\delta_H(v) = m_o(v)$ for every $v \in U$. The existence of such an edge set is straightforward.

Theorem 6.19 (Frank, 1999 [24]). *Let U be a ground-set with a degree-prescription (m_i, m_o) . Let p be a non-negative, integer valued positively crossing supermodular set function on U with $p(\emptyset) = p(U) = 0$. Then there exists an (m_i, m_o) -prescribed edge set H with*

$$\rho_H(X) \geq p(X) \text{ for every } X \subseteq V \tag{6.5}$$

and if and only if

$$m_i(X) \geq p(X) \quad \text{and} \tag{6.6}$$

$$m_o(U - X) \geq p(X) \quad \text{for every } X \subseteq U. \tag{6.7}$$

Theorem 1.34 is an easy consequence: consider a digraph $D = (U + z, A)$ which is (k, ℓ) -edge-connected in U with root node $r_0 \in U$. Let A' denote the set of edges induced by U . For $v \in U$, let $m_o(v) = \delta_A(v, z)$ and $m_i(v) = \delta_A(z, v)$. Let $p(\emptyset) = p(V) = 0$ and let $p(X) = (\gamma(X) - \rho_{A'}(X))^+$ otherwise. It is easy to check that this function is positively crossing supermodular and that the conditions of the theorem are met due to the (k, ℓ) -connectedness in U . The edge set H ensured by the theorem corresponds to the split edges.

Let us now present a generalization of this theorem. The only difference will be that we require a property slightly weaker than positively crossing supermodularity. This is still only a special case of a theorem in the master thesis of T. Király [52, Theorem 2.8]. Our proof follows the same lines as the proof given in [33] for Theorem 6.19. Whereas Theorem 6.19 can be derived from Theorem 1.1, such a deduction does not seem to be possible in our case since we have a skew supermodular-type property.

Theorem 6.20. *Let U be a ground-set with a degree-prescription (m_i, m_o) . Let p be a non-negative, integer valued set function on U with $p(\emptyset) = p(U) = 0$ satisfying the following property. For crossing sets $X, Y \in U$, with $p(X), p(Y) > 0$, either*

$$p(X) + p(Y) \leq p(X \cap Y) + p(X \cup Y) \text{ or} \quad (6.8a)$$

$$p(X) + p(Y) < p(X - Y) + p(Y - X) + m_i(X \cap Y) - m_o(X \cap Y). \quad (6.8b)$$

Then there exists an (m_i, m_o) -prescribed edge set H satisfying (6.5) if and only if (6.6) and (6.7) hold.

Proof. Necessity is obvious as $p(X) \leq \rho_H(X) \leq \min\{m_i(X), m_o(U - X)\}$. For sufficiency, assume for a contradiction that no such H exists. For an (m_i, m_o) -prescribed edge set H , Let $q_H(X) = p(X) - \rho_H(X)$ denote the violation of (6.5) for X and let $\nu_H = \max_{X \subseteq U} q_H(X)$ denote the maximum violation. Let $\mathcal{F}_H := \{X \subset U : q_H(X) = \nu_H\}$ the set of maximally violating sets.¹ As in Section 1.3, assume H is chosen so that ν_H is as small as possible, and subject to this, $|\mathcal{F}_H|$ is as small as possible. As (6.5) does not hold, $\nu_H > 0$, and thus $p(X) > 0$ for every $X \in \mathcal{F}_H$. The next claim is a directed analogue of Claim 5.3.

Claim 6.21. *Let $X, Y \in \mathcal{F}_H$ crossing. Then both $X \cap Y$ and $X \cup Y$ belong to \mathcal{F}_H .*

Proof. If (6.8a) holds for X and Y then $2\nu_H = p(X) + p(Y) - \rho_H(X) - \rho_H(Y) \leq p(X \cup Y) + p(X \cap Y) - \rho_H(X \cup Y) - \rho_H(X \cap Y) \leq 2\nu_H$, hence the claim follows. Assume now (6.8b) holds. Observe that $m_i(X \cap Y) - m_o(X \cap Y) = \rho_H(X \cap Y) - \delta_H(X \cap Y)$. Using this,

$$\begin{aligned} 2\nu_H &= p(X) + p(Y) - \rho_H(X) - \rho_H(Y) < \\ &< p(X - Y) + p(Y - X) + (m_i(X \cap Y) - m_o(X \cap Y)) - \rho_H(X) - \rho_H(Y) \leq \\ &\leq 2\nu_H + \rho_H(X - Y) + \rho_H(Y - X) + (\rho_H(X \cap Y) - \delta_H(X \cap Y)) - \rho_H(X) - \rho_H(Y). \end{aligned}$$

¹It is a difference between the undirected and directed setting that in Section 1.3, \mathcal{F} denoted the set of maximally violating sets minimal for containment.

Finally we get

$$\rho_H(X) + \rho_H(Y) < \rho_H(X - Y) + \rho_H(Y - X) + (\rho_H(X \cap Y) - \delta_H(X \cap Y)),$$

a contradiction to (6.3b). \square

Let K be a minimal member of \mathcal{F} and $L \supseteq K$ be a maximal member. There is an edge $e = uv$ of H with $u, v \in K$ and an $f = xy$ with $x, y \in U - L$ as otherwise K or L would violate (6.6) or (6.7). Let H' be the result of flipping the edges xy and uv , that is, replacing them by uy and xv .

Now $\rho_{H'}(X) \geq \rho_H(X) - 1$ for every $X \subseteq V$ and equality may hold only if $X \cap \{x, y, u, v\}$ is either $\{x, v\}$ or $\{u, y\}$. This condition cannot hold for an $X \in \mathcal{F}$ as it would imply that X and K are crossing. Therefore, $\nu_{H'} \leq \nu_H$ and here equality holds by the minimality of ν_H .

$K \notin \mathcal{F}_{H'}$ as $\rho_{H'}(K) = \rho_H(K) + 1$. So by the minimality of \mathcal{F}_H , there is an $X \in \mathcal{F}_{H'} - \mathcal{F}_H$ with $q_H(X) = \nu_H - 1$. By symmetry we may assume $X \cap \{x, y, u, v\} = \{x, v\}$. $p(X), p(K) > 0$. Again (6.8a) gives a contradiction easily, and if (6.8b) holds, then

$$\begin{aligned} 2\nu_H - 1 &= p(X) + p(K) - \rho_H(X) - \rho_H(K) < \\ &< p(X - K) + p(K - X) + m_i(X \cap K) - m_o(K \cap X) - \rho_H(X) - \rho_H(K) \leq \\ &\leq 2\nu_H - 1 + \rho_H(X - K) + \rho_H(K - X) + \rho_H(X \cap K) - \delta_H(X \cap K) - \rho_H(X) - \rho_H(K). \end{aligned}$$

In the last equation we have used that by the minimal choice of K and $K - X \neq \emptyset$, $q_H(K - X) \leq \nu_H - 1$. This is again a contradiction to (6.3b). \square

We are in the position to derive Lemma 6.13 as an easy consequence.

Proof of Lemma 6.13. Let $F = F_0$. As $F \subseteq \Delta^{out}(r_0)$, it follows that $\mu(X) = \rho_F(X) = \delta_F(s, X)$ for every X . Observe that in this case we only have to guarantee (6.2c) as it implies both (6.2a) and (6.2b).

Let $U = V - u$, and let $D' = (U, A')$ denote the subgraph induced by U . Let us define $p(X)$ the following way. $p(\emptyset) := p(V) := 0$, and for $\emptyset \neq X \neq V$, let

$$p(X) := (\gamma(X) - \rho_{A'}(X) + \mu(X) - t(X))^+ = (\gamma(X) - \rho_{A'-F}(X) - t(X))^+$$

Let $m_o(z) = \delta_D(z, u)$ and $m_i(z) = \delta_D(u, z)$.

Claim 6.22. *The conditions of Theorem 6.20 are satisfied.*

Using this claim Lemma 6.13 follows immediately. Let us split off the edges incident to u according to the edge set A given by the theorem. As u was not special, the edges in F are left unchanged. Let $D_u = (U, A' + H)$ denote the digraph after the splitting. We have to prove that F is globally admissible in D_u . Again it is enough to verify (6.2c), which is a direct consequence of $\rho_H(X) \geq p(X)$.

Proof of Claim 6.22. Consider crossing sets $X, Y \subseteq U$ with $p(X), p(Y) > 0$. Then $t(X) \geq t(X \cup Y)$; furthermore, if X has a seed in $X \cap Y$, then $t(X) = t(X \cap Y)$ and the same holds for exchanging X and Y . Consequently, if $X \cap Y - S \neq \emptyset$ or $X \cap Y$ is special but it contains a seed of X or Y , then $t(X) + t(Y) \geq t(X \cap Y) + t(X \cup Y)$ follows. In this case

$$\begin{aligned} p(X) + p(Y) &= \gamma(X) + \gamma(Y) - t(X) - t(Y) - \rho_{A'-F}(X) - \rho_{A'-F}(Y) \leq \\ &\leq \gamma(X \cup Y) + \gamma(X \cap Y) - t(X \cup Y) - t(X \cap Y) - \\ &\quad - \rho_{A'-F}(X \cup Y) - \rho_{A'-F}(X \cap Y) \leq p(X \cup Y) + p(X \cap Y), \end{aligned}$$

and thus (6.8a) holds. Assume now $X \cap Y$ is special and X has a seed $x \in X - Y$, Y has a seed $y \in Y - X$.

$$\begin{aligned} p(X) + p(Y) &= \gamma(X) + \gamma(Y) - t(X) - t(Y) - \rho_{A'-F}(X) - \rho_{A'-F}(Y) \leq \\ &\leq \gamma(X - Y) + \gamma(Y - X) - t(X) - t(Y) - \\ &\quad - \rho_{A'-F}(X - Y) - \rho_{A'-F}(Y - X) - (\rho_{A'-F}(X \cap Y) - \delta_{A'-F}(X \cap Y)) \end{aligned}$$

As F was insufficient, $|F_t| < \rho_A(t) - \delta_A(t)$ in the original digraph D for every $t \in X \cap Y$, which implies $|F_t| < \rho_{A'}(t) + m_i(t) - \delta_{A'}(t) - m_o(t)$. This gives $m_o(t) - m_i(t) < \rho_{A'-F}(t) - \delta_{A'-F}(t)$, and thus $m_o(X \cap Y) - m_i(X \cap Y) < \rho_{A'-F}(X \cap Y) - \delta_{A'-F}(X \cap Y)$. Now $t(X) = t(X - Y)$ and $t(Y) = t(Y - X)$ because of the seeds x and y , so we get

$$\begin{aligned} p(X) + p(Y) &< \gamma(X - Y) + \gamma(Y - X) - t(X - Y) - t(Y - X) - \\ &\quad - \rho_{A'-F}(X - Y) - \rho_{A'-F}(Y - X) + (m_i(X \cap Y) - m_o(X \cap Y)) \leq \\ &\leq p(X - Y) + p(Y - X) + m_i(X \cap Y) - m_o(X \cap Y). \end{aligned}$$

It is left to verify (6.6) and (6.7). Let $X \subseteq U$. As F was globally admissible in D , $\rho_{A-F}(X) \geq \gamma(X) - t(X)$. Now $\rho_{A-F}(X) = m_i(X) + \rho_{A'-F}(X)$, giving (6.6). On the other hand, $\rho_{A-F}(X + u) \geq \gamma(X + u) - t(X + u) = \gamma(X)$ as $u \notin S$. $\rho_{A-F}(X + u) = m_o(U - X) + \rho_{A'-F}(X)$ and thus $m_o(U - X) \geq \gamma(X) - \rho_{A'-F}(X)$, giving (6.7). \square

\square

6.4 Lemmas

In all claims and lemmas of this sections, F is assumed to be an insufficient globally admissible edge set, if not asserted explicitly otherwise.

Claim 6.23. *Assume $\emptyset \neq Z \subsetneq X \subsetneq V$, $X - Z \subseteq S$ and $\delta_{A-F}(Z, X - Z) = \emptyset$. Then $\rho(Z) < \rho(X) - \delta_F(V - X, X - Z)$ and $\rho_{A-F}(Z) < \rho_{A-F}(X)$.*

Proof. For the first part, $\delta(X - Z) < \rho_{A-F}(X - Z)$ by Claim 6.11 as $X - Z$ is special. Then $\rho(Z) = \rho(X) + \delta(X - Z, Z) - \delta_F(V - X, X - Z) - \delta_{A-F}(V - X, X - Z) < \rho(X) - \delta_F(V - X, X - Z)$ since $\delta(X - Z, Z) - \delta_{A-F}(V - X, X - Z) = \delta(X - Z, Z) - \rho_{A-F}(X - Z) \leq \delta(X - Z) - \rho_{A-F}(X - Z) < 0$ by the previous remark. The second part follows from this using $\rho_F(Z) + \delta_F(V - X, X - Z) \geq \rho_F(X)$. \square

The next lemma describes strong connectivity properties of various tight sets.

Lemma 6.24. (i) Assume X is an out-tight set. If for some $Z \subseteq X$, $\delta_{A-F}(Z, X - Z) = 0$, then Z is out-tight and $\Delta_{D-F}^{\text{out}}(Z) = \Delta_{D-F}^{\text{out}}(X)$. (ii) If X is normal in-tight, $Z \subseteq X$, then $\delta_{A-F}(Z, X - Z) = 0$ implies that $X - Z$ is also normal in-tight and $\Delta_{D-F}^{\text{in}}(X) = \Delta_{D-F}^{\text{in}}(X - Z)$. (iii) If X is μ -tight, and u is a seed of X , then there is an edge $uv \in A - F$ with $v \in X$. (iv) If X is T -tight and v is a sprout of X , then there is an edge $uv \in A - F$ with $u \in X$.

Proof. (i) $\delta_{A-F}(X) = \ell$ and $\delta_{A-F}(Z) \geq \ell$. Thus if $\delta_{A-F}(Z, X - Z) = 0$ then all edges in $A - F$ leaving Z must leave X as well, and this is what we wanted to prove.

(ii) Assume first $X - Z - S \neq \emptyset$. $\rho_{A-F}(X) = k$, $\rho_{A-F}(X - Z) \geq k$, and the claim follows as in the first part.

Assume now $X - Z$ is special. By Claim 6.23, $\rho_{A-F}(Z) < \rho_{A-F}(X) = k$, a contradiction as X was not special, and thus neither is Z .

(iii) $\rho(X) = k + \delta_F(V - X - S, X - u)$. If all edges in X outgoing from u are in F , then we can use Claim 6.23 for $Z = \{u\}$, and thus $k = \rho(u) < k + \delta_F(V - X - S, X - u) - \delta_F(V - X, X - u) \leq k$, a contradiction.

(iv) $\rho(X) = k + T(X) = k + \delta_F(V - X, v)$. If all edges in X entering v are in F , then Claim 6.23 can be applied for $Z = X - v$. Thus $k \leq \rho(X - v) < k + T(X) - \delta_F(V - X, v) = k$, a contradiction again. \square

Claim 6.25. For sets $\emptyset \neq Z \subseteq X$, $X - Z \subseteq S$, if X has a seed $u \in Z$ then $t(X) = t(Z)$.

Proof. As $X - Z \subseteq S$, for any $x \in Z$, $\delta_F(V - Z - S, x) = \delta_F(V - X - S, x)$. u is the node in X minimizing $\delta(V - X - S, x)$, and thus the claim follows. \square

In the next lemma, we show some configurations of tight sets which may not exist for an insufficient globally admissible F .

Lemma 6.26. There exists no $X \subseteq V$ with the following properties: $|X| \geq 2$, X is in-tight and (i) $X - S \neq \emptyset$ and there is a subpartition $\mathcal{Y} = \{Y_1, \dots, Y_m\}$ of X so that $X - S \subseteq \cup \mathcal{Y}$ and each Y_i is out-tight and proper subset of X or (ii) X is μ -tight and there is an out-tight $Y \subsetneq X$ containing a seed u of X ; (iii) X is T -tight and there is an out-tight $Y \subsetneq X$ not containing a sprout z of X .

Proof. (i) We may assume that there is no special Y_i as leaving such members from \mathcal{Y} the conditions still hold. Thus $\rho_{A-F}(Y_i) \geq k$ for each i and $\delta_{A-F}(Y_i) = \ell$ as they are out-tight sets. Let $X_0 = X - \cup \mathcal{Y}$. As X_0 is special, Claim 6.11 implies $\rho_{A-F}(X_0) - \delta_{A-F}(X_0) > \delta_F(X_0)$ whenever $X_0 \neq \emptyset$. Now $\rho_{A-F}(X) = k$, $\delta_{A-F}(X) \geq \ell$, and thus

$$\begin{aligned} k - \ell &\geq \rho_{A-F}(X) - \delta_{A-F}(X) = \\ &= (\rho_{A-F}(X_0) - \delta_{A-F}(X_0)) + \sum_{i=1}^m (\rho_{A-F}(Y_i) - \delta_{A-F}(Y_i)) \geq \delta_F(X_0) + m(k - \ell), \end{aligned}$$

a contradiction, since either $X_0 \neq \emptyset$ and thus the last inequality is strict, or $m \geq 2$ as we did not allow $\mathcal{Y} = \{X\}$.

(ii) Let u denote a seed of X as in the conditions. $t(X) = t(Y)$ by Claim 6.25 ($X - Y \subseteq S$ holds since X is special). $\delta(Y) = \ell + \delta_F(Y)$ as Y is out-tight. Claim 6.11 gives $\rho(X - Y) - \delta(X - Y) > \rho_F(X - Y)$. Similarly to the previous case,

$$\begin{aligned} k + \mu(X) - t(X) - \ell - \delta_F(X) &\geq \rho(X) - \delta(X) = \rho(X - Y) - \delta(X - Y) + \\ &+ \rho(Y) - \delta(Y) > \rho_F(X - Y) + k + \mu(Y) - t(Y) - \ell - \delta_F(Y). \end{aligned}$$

This gives $\delta_F(Y) - \delta_F(X) + \mu(X) - \mu(Y) > \rho_F(X - Y)$. Using $\delta_F(Y) \leq \delta_F(X) + \delta_F(Y, X - Y)$ and $\mu(X) = \mu(Y) + \delta_F(V - X - S, X - Y)$, one gets $\delta_F(Y, X - Y) + \delta_F(V - X - S, X - Y) > \rho_F(X - Y)$, clearly a contradiction.

(iii) As in the previous two cases,

$$\begin{aligned} k + T(X) - \ell - \delta_F(X) &\geq \rho(X) - \delta(X) = \rho(X - Y) - \delta(X - Y) + \\ &+ \rho(Y) - \delta(Y) > \rho_F(X - Y) + k - \ell - \delta_F(Y). \end{aligned}$$

Thus $\delta_F(Y) - \delta_F(X) + T(X) > \rho_F(X - Y)$. As $\delta_F(Y) \leq \delta_F(X) + \delta_F(Y, X - Y)$ and $T(X) = \delta_F(V - X, z)$, we have $\delta_F(Y, X - Y) + \delta_F(V - X, z) > \rho_F(X - Y)$, a contradiction again. \square

Claim 6.27. (a) If $X \cap Y$ is special, then $\rho(X) + \rho(Y) > \rho(X - Y) + \rho(Y - X) + \delta_F(V - X, X \cap Y) + \delta_F(V - Y, X \cap Y)$.

(b) If Y is normal tight, $Y - X - S \neq \emptyset$, $r_0 \notin X \cap Y$, then $\rho(Y) \leq \rho(Y - X) + \delta_F(V - Y, X \cap Y)$.

Proof. (a) By (6.3b), it is enough to prove that $(\rho(X \cap Y) - \delta(X \cap Y)) + \bar{d}(X, Y) > \delta_F(V - X, X \cap Y) + \delta_F(V - Y, X \cap Y)$. By Claim 6.11, $\rho_F(X \cap Y) < \rho(X \cap Y) - \delta(X \cap Y)$ and obviously, $\delta_F(V - X - Y, X \cap Y) \leq \bar{d}(X, Y)$. These together imply the claim.

(b) Since $Y - X$ is not special, $\rho(Y - X) \geq \gamma(Y - X) + \rho_F(Y - X)$ and $\gamma(Y - X) = \gamma(Y)$ as $r_0 \notin X \cap Y$. Using these,

$$\begin{aligned} \rho(Y) &= \gamma(Y) + \rho_F(Y) = \gamma(Y) + \delta_F(V - Y, Y - X) + \delta_F(V - Y, X \cap Y) \leq \\ &\leq \gamma(Y - X) + \rho_F(Y - X) + \delta_F(V - Y, X \cap Y) \leq \rho(Y - X) + \delta_F(V - Y, X \cap Y). \end{aligned}$$

\square

We are almost ready to prove Lemma 6.14. The following lemma is slightly weaker, but will easily imply it.

Lemma 6.28. *If F' is globally admissible and there exists at least one special tight set, then any maximal globally admissible set $F \supseteq F'$ is sufficient.*

Proof. Let F be a maximal globally admissible set containing F' . Clearly, the tight sets for F are also tight for F' . We show that if F is insufficient, then no special tight set may exist.

First we show that no T -tight set exists. Indeed, assume X is minimal T -tight; let z be a sprout. By Lemma 6.24(iv), there is an edge $uz \in A - F$ with $u \in X$. By Claim 6.6, uz must enter a tight set Y which is either normal or T -tight with sprout z . Case (c) is excluded since u is special.

First assume Y is normal. If $V - Y \subseteq X$ then we have a contradiction by Lemma 6.26(iii) as $V - Y$ is an out-tight set satisfying the conditions. $Y \subset X$ is impossible as it would give $Y \subseteq S$. Thus X and Y are crossing.

$$\rho(X) = k + T(X) \leq \rho(X - Y) + \delta_F(V - X, X \cap Y) \quad (6.9)$$

as $z \in X \cap Y$ and $\rho(X - Y) \geq \gamma(X - Y) = k$. Using both Claim 6.27(b) and (a) we get a contradiction unless F is sufficient.

If Y is a T -tight set, by the minimality of X , X and Y are crossing. (6.9) holds again and also $\rho(Y) = k + T(Y) \leq \rho(Y - X) + \delta_F(V - Y, X \cap Y)$ as $z \in X \cap Y$ is also a sprout of Y . A contradiction again.

Next, assume X is minimal μ -tight, and let u be a seed. By Lemma 6.24(iii), we have a $uv \in A - F$ with $v \in X$ blocked by a tight set Y . We have seen already that no T -tight sets exist. Neither may Y be μ -tight since u is special. Thus Y should be normal. Again $V - Y \subseteq X$ would contradict Lemma 6.26(ii) and $Y \subset X$ is impossible, and thus X and Y should be crossing. Using Claim 6.25 for X and $Z = X - Y$, $t(X - Y) = t(X)$. Thus

$$\begin{aligned} \rho(X) &= k + \mu(X) - t(X) = k + \delta_F(V - S - X, X) - t(X - Y) = \\ &k + \delta_F(V - S - X, X - Y) - t(X - Y) + \delta_F(V - S - X, X \cap Y) \leq \\ &\leq \rho(X - Y) + \delta_F(V - X, X \cap Y). \end{aligned}$$

Using again Claim 6.27(b) and (a) gives a contradiction. \square

Lemma 6.29. *Assume F is a maximal, insufficient globally admissible set of edges. If X and Y are crossing tight sets, then $X \cup Y$ and $X \cap Y$ are tight as well. If X or Y blocks an edge $uv \in A - F$, then either $X \cup Y$ or $X \cap Y$ blocks uv as well.*

Proof. By Lemma 6.28, we know that both X and Y are normal tight. Assume first that $(X \cap Y) - S \neq \emptyset$. From (6.3a) and (6.4a) we have:

$$\begin{aligned} \rho_{A-F}(X) + \rho_{A-F}(Y) &= \gamma(X) + \gamma(Y) = \gamma(X \cap Y) + \gamma(X \cup Y) \leq \\ &\leq \rho_{A-F}(X \cap Y) + \rho_{A-F}(X \cup Y) \leq \rho_{A-F}(X) + \rho_{A-F}(Y), \end{aligned}$$

implying that both $X \cap Y$ and $X \cup Y$ are tight and $d_{A-F}(X, Y) = 0$. The second part of the claim follows as both of them are normal.

We show that $X \cap Y \subseteq S$ is impossible. $X - Y$ and $Y - X$ are both non-special sets, and thus Claim 6.27(b) applies for Y and also for X by exchanging the role of X and Y . Claim 6.27(a) leads to a contradiction again. \square

An easy consequence of Lemma 6.29 is the following:

Claim 6.30. *If F is maximal insufficient globally admissible and $uv \in A - F$, either there is a unique minimal in-tight set B_{uv}^{in} blocking uv or a unique minimal out-tight B_{uv}^{out} blocking uv . If $u, v \in X$ for an in- or out-tight set X , then $B_{uv}^{in} \subseteq X$ or $B_{uv}^{out} \subseteq X$.*

Proof. By Lemma 6.29, for every edge $uv \in A - F$ there is a unique minimal B_1 and a unique maximal B_2 in-tight set entered by uv . If $r_0 \notin B_1$ then B_1 is in-tight and thus $B_{uv}^{in} = B_1$, if $r_0 \in B_1$ then $B_{uv}^{out} = V - B_2$. (Note that both sets may exist). The second part also follows by Lemma 6.29. \square

Now we are ready to prove Lemmas 6.14 and 6.15.

Proof of Lemma 6.14. By Lemma 6.28, the only case left is if X is normal tight with $r_0 \notin X$, $|X - S| = 1$. Let $X - S = \{u\}$. If there is no edge in $A - F$ from u to $X - u$, then by Lemma 6.24, $X - u$ is normal in-tight, a contradiction to $X - u \subseteq S$. Thus there exists an edge $uv \in A - F$ with $v \in X$. Let $Y = B_{uv}^{in}$ or $Y = B_{uv}^{out}$ as in Claim 6.30. In the first case $Y \subseteq S$ contradicting that it is a tight set and every tight set is normal. In the second case, X and $\mathcal{Y} = \{Y\}$ satisfy the conditions of Lemma 6.26(i), a contradiction again. \square

Proof of Lemma 6.15. For a contradiction, assume F is insufficient. Let K denote the set of in-tight singletons and L the set of out-tight singletons.

Claim 6.31. $K \cap L = \emptyset$.

Proof. Let $u \in K \cap L$. Trivially, $u \neq r_0$. As a singleton tight set cannot be special, $\rho(u) = k$ and $\delta(u) \geq k$. However, the out-tightness of $\{u\}$ implies $\delta_{A-F}(u) = \ell$, and thus $\delta_F(u) > 0$, a contradiction. \square

Claim 6.32. *If an edge $f = xy \in A - F$ is blocked by an in-tight set, then $B_{xy}^{in} = \{y\}$. If it is blocked by an out-tight set, then $B_{xy}^{out} = \{x\}$.*

Proof. Consider a minimal in-tight or out-tight set X for some edge $f = xy \in A - F$ which is not a singleton. By Lemma 6.24(i) or (ii) and the minimality of X , X is strongly connected in $A - F$. We show that either $X \subseteq K$ or $X \subseteq L$. Consider an edge $uv \in A - F$ with $u, v \in X$, guaranteed by the strong connectivity. By Claim 6.30, either uv enters a minimal in-tight or

leaves a minimal out-tight Y with $Y \subseteq X$. By the minimal choice of X , Y is a singleton: $Y = \{u\} \in L$ or $Y = \{v\} \in K$. Thus either $X \cap K \neq \emptyset$ or $X \cap L \neq \emptyset$.

Assume first $X \cap K \neq \emptyset$ and let $Z = X \cap K$. If $X - Z \neq \emptyset$, then by the strongly connectedness there is an edge $uv \in A - F$ with $u \in Z$ and $v \in X - Z$ blocked by a minimal in- or out-tight set Y . Again, Y is a singleton and either $Y = \{u\} \in L$ or $Y = \{v\} \in K$. Both cases are impossible since $u \in X \cap K$, and $v \in X - K$. Thus we may conclude $X \subseteq K$.

Next, consider $X \cap L \neq \emptyset$ and let $Z = X \cap L$. If $X - Z \neq \emptyset$, then an edge $uv \in A - F$ with $u \in X - Z$, $v \in Z$ gives the contradiction as above. Thus $X \subseteq L$ follows.

X was either in- or out-tight. If $X = B_{xy}^{out}$ is out-tight, then $X \subseteq L$ is excluded as it would give $B_{xy}^{out} = \{x\}$. Thus $X \subseteq K$. As $K \cap S = \emptyset$, for each $u \in X$, $\rho(u) = k$, $\delta(u) \geq k$. By the assumption that all edges in F have tail in $S + r_0$, $\delta_F(X) = 0$ and thus $\delta(X) = \ell$. Now

$$k - \ell \leq \rho(X) - \delta(X) = \sum_{u \in X} (\rho(u) - \delta(u)) \leq 0,$$

giving a contradiction.

If $X = B_{xy}^{in}$ is in-tight, then $X \subseteq K$ is excluded since it would give $B_{xy}^{in} = \{y\}$. Thus $X \subseteq L$. $X - S \neq \emptyset$ as all tight sets are normal by Lemma 6.28, and thus the conditions of Lemma 6.26(i) apply with \mathcal{Y} being the partition of X into singletons. \square

$r_0 \notin K$ implies $K \neq V$. Also $K \neq \emptyset$ as by Claim 6.32, all edges in $A - F$ leaving r_0 should enter members of K . As $\rho_{A-F}(V - K) \geq \ell$, there is an edge $uv \in A - F$ leaving K . This cannot be blocked by neither an in-tight nor an out-tight singleton. \square

6.5 Further remarks

6.5.1 Matroid property of locally admissible sets

First, we describe the structure of the locally admissible edge sets at a given special node z . We prove

Theorem 6.33. *The set system $M_z = \{F : F \text{ is locally admissible at } z\}$ is a matroid.*

This together with Theorem 6.1 gives a straightforward way for finding a sufficient locally admissible edge set. By Theorem 6.1, we know that special nodes exist and one of them has a sufficient locally admissible set. We check the special nodes one-by-one, and at each special node z we greedily choose a maximal locally admissible edge set. Note that this can be done easily as we just need to take care of the (k, ℓ) -edge-connectedness in $V - z$ which can be checked by flow computations. Theorem 6.33 ensures that if z admits a sufficient global admissible edge set, we can find it this way.

Proof of Theorem 6.33. The only nontrivial property we have to check is that if $|F| < |F'|$ and both $F, F' \in M_z$ then there is an edge $uz \in F' - F$ so that $F + uz$ is locally admissible as well. For a contradiction, assume this does not hold.


A set X will now be called **tight** at z for F if $z \in X$, $X \neq \{z\}$ and it satisfies (6.1) with equality. (Actually this notion coincides with the tight sets containing z when we consider F as a globally admissible set of edges). Note that since $|F'| \leq k - \delta(z)$ by definition and $|F| < |F'|$, $|F|$ is insufficient.

Claim 6.34. *If X and Y are crossing tight sets at z for F then $X \cap Y$ and $X \cup Y$ are also tight.*

Proof. If $X \cap Y \neq \{z\}$, then (6.1) also holds for $X \cap Y$ and $X \cup Y$ and thus the claim follows by the submodularity of the function ρ_{A-F} . We show that $X \cap Y = \{z\}$ is impossible. Indeed, by (6.3b) we would have $\gamma(X) + \gamma(Y) = \rho_{A-F}(X) + \rho_{A-F}(Y) \geq \rho_{A-F}(X - Y) + \rho_{A-F}(Y - X) + \rho_{A-F}(z) - \delta_{A-F}(z) > \rho_{A-F}(X - Y) + \rho_{A-F}(Y - X) \geq \gamma(X - Y) + \gamma(Y - X)$ as F was insufficient, a contradiction to (6.4b). \square

Thus for each edge $uz \in F' - F$ there is a unique minimal tight set X_{uz} at z for F entered by uz . For different $uz, wz \in F' - F$, X_{uz} and X_{wz} cannot be crossing as $X_{uz} \cap X_{wz}$ would also be tight contradicting their minimality. Thus $X_{uz} \cup X_{wz} = V$. Let $\mathcal{T} = \{V - X_{uz} : uz \in F' - F\}$. \mathcal{T} forms a subpartition of $V - z$ so that for each $uz \in F' - F$, u is contained in some member of \mathcal{T} . For each $Y \in \mathcal{T}$, $\delta(Y) = \gamma(V - Y) + \delta_F(Y)$. As F' is locally admissible, $\delta_{F'}(Y) \leq \delta(Y) - \gamma(V - Y) = \delta_F(Y)$, and thus $\delta_{F'-F}(Y) \leq \delta_{F-F'}(Y)$. Summing up for all $Y \in \mathcal{T}$ we get $|F' - F| = \sum_{Y \in \mathcal{T}} \delta_{F'-F}(Y) \leq \sum_{Y \in \mathcal{T}} \delta_{F-F'}(Y) \leq |F - F'|$, contradicting $|F| < |F'|$. \square

6.5.2 Example of an insufficient maximal globally admissible set



An example for an insufficient maximal globally admissible set is shown on the figure for $k = 4$, $\ell = 2$. D is minimally $(4, 2)$ -edge-connected. It contains two special nodes u and t with in-degree 4 and out-degree 2. Both of them have a sufficient locally admissible edge set: for both u and t the two edges coming from w are sufficient locally admissible. However, if we

consider F consisting of one wu and one wt edge (the thick edges), F is maximal as the following sets block every edge entering u and t : $\{u\}$, $\{t\}$, $\{w\}$ are out-tight and $\{u, t, v, w\}$ is in-tight. However, F is insufficient.

The proof of the case $\ell = k - 1$ by Frank and Király [33] used an argument similar to the proof of Lemma 6.15. One might wonder why the much simpler argument cannot be applied in the general case to prove that every maximal globally admissible set is sufficient (which is, in fact, false). A possible explanation is that Claim 6.31 fails to hold unless F satisfies the condition in Lemma 6.15: in this example the singleton set $\{w\}$ is both in- and out-tight.

Bibliography

- [1] R. Ahuja, T. Magnanti, and J. Orlin. *Network flows. Theory, algorithms and applications*. Prentice-Hall, New York, 1993. [84](#)
- [2] J. Bang-Jensen, H. N. Gabow, T. Jordán, and Z. Szigeti. Edge-connectivity augmentation with partition constraints. *SIAM J. Discrete Math.*, 12(2):160–207, 1999. [14](#), [15](#), [27](#), [30](#)
- [3] J. Bang-Jensen and B. Jackson. Augmenting hypergraphs by edges of size two. *Math. Program.*, 84(3):467–481, 1999. [14](#)
- [4] A. A. Benczúr. Pushdown-reduce: an algorithm for connectivity augmentation and poset covering problems. *Discrete Appl. Math.*, 129(2-3):233–262, 2003. [10](#), [23](#), [73](#)
- [5] A. A. Benczúr, J. Förster, and Z. Király. Dilworth’s theorem and its application for path systems of a cycle - implementation and analysis. In *ESA '99: Proceedings of the 7th Annual European Symposium on Algorithms*, pages 498–509, London, UK, 1999. Springer-Verlag. [10](#)
- [6] A. A. Benczúr and A. Frank. Covering symmetric supermodular functions by graphs. *Mathematical Programming*, 84(3):483–503, 1999. [13](#), [14](#), [27](#)
- [7] K. Bérczi and Y. Kobayashi. An algorithm for $(n - 3)$ -connectivity augmentation problem: Jump system approach. Technical Report METR 2009-12, Department of Mathematical Engineering, University of Tokyo, April 2009. [6](#), [10](#)
- [8] K. Bérczi and L. A. Végh. Restricted b -matchings in degree-bounded graphs. Technical Report TR-2009-12, Egerváry Research Group, Budapest, 2009. www.cs.elte.hu/egres. [6](#), [23](#)
- [9] A. Bernáth. *Edge-connectivity augmentation of graphs and hypergraphs*. PhD thesis, Eötvös University, Budapest, 2009. [15](#), [26](#)
- [10] A. Bernáth. A simple proof of a theorem of Benczúr and Frank. Technical Report TR-2009-02, Egerváry Research Group, Budapest, 2009. www.cs.elte.hu/egres. [27](#)

- [11] A. Bernáth, R. Grappe, and Z. Szigeti. Partition constrained covering of a symmetric crossing supermodular function by a graph. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA10)*, 2010. (to appear). [15](#), [112](#)
- [12] A. Bernáth and T. Király. A new approach to splitting-off. In *Proceedings of the 13th International Conference on Integer Programming and Combinatorial Optimization, IPCO 2008*, volume 5035 of *Lecture Notes in Computer Science*, pages 401–415. Springer, 2008. (See the full version in the EGRES technical report "TR-2008-02" at www.cs.elte.hu/egres.)
[27](#)
- [13] J. Cheriyan and R. Thurimella. Fast algorithms for k-shredders and k-node connectivity augmentation. *Journal of Algorithms*, 33(1):15–50, 1999. [10](#)
- [14] G. Cornuéjols and W. Pulleyblank. A matching problem with side conditions. *Discrete Math.*, 29:135–139, 1980. [6](#)
- [15] B. Cosh. *Vertex Splitting and Connectivity Augmentation in Hypergraphs*. PhD thesis, University of London, 2000. [27](#)
- [16] J. Edmonds. Existence of k-edge connected ordinary graphs with prescribed degrees. *J. Res. Natl. Bur. Stand.*, 68B:73–74, 1964. [27](#)
- [17] J. Edmonds. Edge disjoint branchings. In B. Rustin, editor, *Combinatorial Algorithms*, pages 91–96. Academic Press, New York, 1973. [17](#)
- [18] S. Enni. A 1-(S, T)-edge-connectivity augmentation algorithm. *Mathematical Programming*, 84, 1999. [10](#)
- [19] K. Eswaran and R. Tarjan. Augmentation problems. *SIAM J. Computing*, 5(4):653–665, 1976. [9](#), [10](#)
- [20] T. Fleiner. Covering a symmetric poset by symmetric chains. *Combinatorica*, 17(3):339–344, 1997. [22](#), [53](#)
- [21] A. Frank. *Combinatorial algorithms, algorithmic proofs (in Hungarian)*. PhD thesis, Eötvös University, Budapest, 1976. [23](#)
- [22] A. Frank. On the orientation of graphs. *J. Comb. Theory, Ser. B.*, 28(3):251–261, 1980. [17](#), [18](#)
- [23] A. Frank. Augmenting graphs to meet edge-connectivity requirements. *SIAM J. Discrete Math.*, 5(1):25–53, 1992. [1](#), [8](#), [9](#), [11](#), [12](#), [13](#), [62](#)

- [24] A. Frank. Connectivity augmentation problems in network design. In J. Birge and K. Murty, editors, *Mathematical Programming: State of the Art*, pages 34–63. The University of Michigan, 1999. [17](#), [18](#), [121](#)
- [25] A. Frank. Finding minimum generators of path systems. *J. Comb. Theory Ser. B*, 75(2):237–244, 1999. [10](#)
- [26] A. Frank. An algorithm to increase the node-connectivity of a digraph. In *Proceedings of the 3rd Hungarian-Japanese Symposium on Discrete Mathematics and Its Applications*, pages 378–387, Tokyo, Japan, 2003. [10](#)
- [27] A. Frank. Restricted t -matchings in bipartite graphs. *Discrete Appl. Math.*, 131(2):337–346, 2003. [6](#)
- [28] A. Frank. Edge-connection of graphs, digraphs, and hypergraphs. In E. Györi, G. Katona, and L. Lovász, editors, *More sets, graphs and numbers*, *Bolyai Mathematical Society Math. Studies*, volume 5, pages 93–142. Springer Verlag, 2006. [31](#)
- [29] A. Frank. Rooted k -connections in digraphs. *Discrete Appl. Math.*, 157(6):1242–1254, 2009. [63](#)
- [30] A. Frank and T. Jordán. Presented at the Workshop on Network Design, Budapest, October 1994. [21](#)
- [31] A. Frank and T. Jordán. Minimal edge-coverings of pairs of sets. *J. Comb. Theory Ser. B*, 65(1):73–110, 1995. [1](#), [2](#), [4](#), [9](#), [39](#), [62](#)
- [32] A. Frank and T. Jordán. Directed vertex-connectivity augmentation. *Math. Prog.*, 84:537–553, 1999. [10](#)
- [33] A. Frank and Z. Király. Graph orientations with edge-connection and parity constraints. *Combinatorica*, 22(1):47–70, 2002. [18](#), [19](#), [27](#), [32](#), [122](#), [131](#)
- [34] A. Frank and L. Szegő. Constructive characterizations for packing and covering with trees. *Discrete Appl. Math.*, 131(2):347–371, 2003. [18](#), [31](#)
- [35] A. Frank and E. Tardos. Generalized polymatroids and submodular flows. *Math. Program.*, 42(3):489–563, 1988. [63](#)
- [36] A. Frank and L. A. Végh. An algorithm to increase the node-connectivity of a digraph by one. *Discrete Optimization*, 5:677–684, 2008. [33](#), [35](#), [61](#), [139](#), [141](#)
- [37] D. S. Franzblau and D. J. Kleitman. An algorithm for covering polygons with rectangles. *Inf. Control*, 63(3):164–189, 1986. [10](#)

- [38] E. Györi. A min-max theorem on intervals. *J. Comb. Theory Ser. B*, 37:1–9, 1984. [5](#)
- [39] S. L. Hakimi. On realizability of a set of integers as degrees of the vertices of a linear graph. I. *Journal of the Society for Industrial and Applied Mathematics*, 10(3):496–506, 1962. [27](#)
- [40] D. Hartvigsen. *Extensions of matching theory*. PhD thesis, Carnegie-Mellon University, 1984. [6](#)
- [41] D. Hartvigsen. The square-free 2-factor problem in bipartite graphs. In *Proceedings of the 7th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pages 234–241, Berlin, 1999. Springer-Verlag. [6](#)
- [42] D. Hartvigsen. Finding maximum square-free 2-matchings in bipartite graphs. *J. Comb. Theory Ser. B*, 96(5):693–705, 2006. [6](#)
- [43] J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matching in bipartite graphs. *SIAM J. Comp.*, 2:225–231, 1973. [84](#)
- [44] T.-S. Hsu. On four-connecting a triconnected graph. *J. Algorithms*, 35(2):202–234, 2000. [10](#)
- [45] T. Ishii and H. Nagamochi. On the minimum augmentation of an l -connected graph to a k -connected graph. In *SWAT '00: Proceedings of the 7th Scandinavian Workshop on Algorithm Theory*, 2000. [11](#)
- [46] B. Jackson and T. Jordán. A near optimal algorithm for vertex connectivity augmentation. In *ISAAC '00: Proceedings of the 11th International Conference on Algorithms and Computation*, pages 312–325, London, UK, 2000. Springer-Verlag. [11](#)
- [47] B. Jackson and T. Jordán. Independence free graphs and vertex connectivity augmentation. *J. Comb. Theory Ser. B*, 94(1):31–77, 2005. [10](#)
- [48] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001. [113](#)
- [49] T. Jordán. On the optimal vertex-connectivity augmentation. *J. Comb. Theory Ser. B*, 63(1):8–20, 1995. [10](#), [48](#), [60](#)
- [50] T. Jordán. A note on the vertex-connectivity augmentation problem. *J. Comb. Theory, Ser. B*, 71(2):294–301, 1997. [10](#), [48](#), [60](#)
- [51] T. Jordan. Two NP-complete augmentation problems. Technical report, 1997. [113](#)
- [52] T. Király. The splitting off operation and its applications (in hungarian). Master’s thesis, Eötvös University, Budapest, 1999. available online at <http://www.cs.elte.hu/>. [122](#)

- [53] Z. Király. C_4 -free 2-factors in bipartite graphs. Technical Report TR-2001-13, Egerváry Research Group, Budapest, 2001. www.cs.elte.hu/egres. 6
- [54] Z. Király, B. Cosh, and B. Jackson. Local edge-connectivity augmentation in hypergraphs is np-complete. Technical Report TR-2009-06, Egerváry Research Group, Budapest, 2009. www.cs.elte.hu/egres. 13, 26
- [55] D. E. Knuth. Irredundant intervals. *J. Exp. Algorithmics*, 1:1, 1996. 10
- [56] E. R. Kovács and L. A. Végh. The constructive characterization of (k, ℓ) -edge-connected digraphs. *Combinatorica*. (accepted); available as EGRES Tech. Report TR-2008-14 at <http://www.cs.elte.hu/egres>. 33, 115, 139, 141
- [57] E. R. Kovács and L. A. Végh. Constructive characterization theorems in combinatorial optimization. In *RIMS Kôkyûroku Bessatsu*, 2009. (to appear). 16
- [58] L. C. Lau and C. K. Yung. Efficient edge splitting and constrained edge splitting. manuscript, 2009. 29
- [59] G. Liberman and Z. Nutov. On shredders and vertex connectivity augmentation. *J. of Discrete Algorithms*, 5(1):91–101, 2007. 10, 48
- [60] L. Lovász. *Combinatorial Problems and Exercises*. Akadémiai Kiadó - North Holland, Budapest, 1979. 13, 16
- [61] L. Lovász and M. D. Plummer. *Matching theory*, volume 121 of *North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam, 1986. *Annals of Discrete Mathematics*, 29. 7
- [62] A. Lubiw. A weighted min-max relation for intervals. *J. Comb. Theory Ser. A*, 53(2):151–172, 1990. 10
- [63] W. Mader. A reduction method for edge-connectivity in graphs. *Annals of discrete Math*, 3:145–164, 1978. 12, 16
- [64] W. Mader. Konstruktion aller n -fach kantenzusammenhängenden digraphen. *Europ. J. Combinatorics*, 3:63–67, 1982. 16, 17
- [65] M. Makai. On maximum cost $K_{t,t}$ -free t -matchings of bipartite graphs. *SIAM J. Discrete Math.*, 21(2):349–360, 2007. 7
- [66] C. S. J. A. Nash-Williams. On orientations, connectivity and odd-vertex-pairings in finite graphs. *Canad. J. Math.* 12, pages 555–567, 1960. 13, 16

- [67] Z. Nutov. Approximating connectivity augmentation problems. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 176–185, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics. 27
- [68] Z. Nutov. Approximating rooted connectivity augmentation problems. *Algorithmica*, 44(3):213–231, 2006. 63
- [69] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003. 3, 36, 39, 60, 62
- [70] J. Szabó. Jump systems and matroid parity (in hungarian). Master’s thesis, Eötvös University, Budapest, 2002. Available online at <http://www.cs.elte.hu>. 62
- [71] W. T. Tutte. On the problem of decomposing a graph into n connected factors. *J. London Math. Soc.*, pages 1–36, 1961. 17
- [72] L. A. Végh. Directed connectivity augmentation (in hungarian). Master’s thesis, Eötvös University, Budapest, 2004. available online at <http://www.cs.elte.hu/>. 42
- [73] L. A. Végh. Augmenting undirected node-connectivity by one. Technical Report TR-2009-10, Egerváry Research Group, Budapest, 2009. <http://www.cs.elte.hu/egres>. 33, 139, 141
- [74] L. A. Végh and A. A. Benczúr. Primal-dual approach for directed vertex connectivity augmentation and generalizations. *ACM Transactions on Algorithms*, 4(2), 2008. 33, 69, 139, 141
- [75] T. Watanabe and A. Nakamura. Edge-connectivity augmentation problems. *J. Comput. Syst. Sci.*, 35(1):96–144, 1987. 1, 9, 11
- [76] T. Watanabe and A. Nakamura. A minimum 3-connectivity augmentation of a graph. *J. Comput. Syst. Sci.*, 46(1):91–128, 1993. 10
- [77] H. Whitney. On the classification of graphs. *Am. J. Math.*, 55:236–244, 1933. 15
- [78] D. P. Williamson, M. X. Goemans, M. Mihail, and V. V. Vaziran. A primal-dual approximation algorithm for generalized steiner network problems. *Combinatorica*, 15:708–717, 1995. 113

Abstract

The main subject of the thesis is connectivity augmentation: we would like to make a given graph k -connected by adding a minimum number of new edges. There are four basic problems in this field, since one might consider both edge- and node-connectivity augmentation in both graphs and digraphs. The thesis wishes to contribute to three out of these four problems: directed- and undirected node-connectivity and undirected edge-connectivity augmentation. Although directed edge-connectivity augmentation is not being considered, the last chapter is devoted to a constructive characterization result related to directed edge-connectivity. Let us summarize the main results of the thesis.

- *We present a min-max formula and a combinatorial polynomial time algorithm for augmenting undirected node-connectivity by one.* The complexity status of undirected node-connectivity augmentation of arbitrary graphs is still open; already the special case of augmenting by one has attracted considerable attention. The formula proved in Chapter 3 was conjectured by Frank and Jordán in 1994.
- *We present the first combinatorial polynomial time algorithm for directed node-connectivity augmentation.* For this problem, Frank and Jordán gave a min-max formula in 1995; however, it remained an open problem to develop a combinatorial algorithm. We present two, completely different combinatorial algorithms. Chapter 2 contains one for the special case of augmenting connectivity by one (a joint work with András Frank), and Chapter 4 presents another for augmenting the connectivity of arbitrary digraphs (a joint work with András Benczúr Jr.). The latter result also gives a new, algorithmic proof of the general theorem of Frank and Jordán on covering positively crossing supermodular functions on set pairs.
- *We establish a constructive characterization of (k, ℓ) -edge-connected digraphs.* This result of Chapter 6, a joint work with Erika Renáta Kovács, settles a conjecture of Frank from 2003. The theorem gives a common generalization of a number of previously known characterizations, and naturally fits into the framework defined by splitting off and orientation theorems.
- *We present partial results concerning partition constrained undirected local edge-connectivity augmentation.* In Chapter 5, we discuss some classical results concerning undirected edge-connectivity augmentation in a unified framework, based on the technique of edge-flippings. For the partition constrained problem we formulate a conjecture and give a partial proof.

Most results are based on the papers [36], [74], [73] and [56], except for Chapter 5, which contains unpublished results.

Összefoglalás

Az értekezés fő témája az összefüggőség-növelés: egy adott gráfot szeretnénk minimális számú él hozzávételével k -szorosán összefüggővé tenni. Ez négy alapkérdést foglal magában, mivel él- és pontösszefüggőség növelése is felvethető mind irányított, mind irányítatlan gráfokban. Az értekezésben ezen alapproblémák közül hárommal foglalkozunk: az irányított és irányítatlan pontösszefüggőség, valamint az irányítatlan élösszefüggőség növelésével. Irányított élösszefüggőség-növelésről ugyan nem esik szó, viszont az utolsó fejezetben ezzel az összefüggőség-fogalommal kapcsolatban adunk egy konstruktív karakterizációs eredményt. Az értekezés fő eredményei a következők.

- *Megadunk egy min-max formulát és egy kombinatorikus polinomiális algoritmust az irányítatlan pontösszefüggőség eggyel való növelésére.* Tetszőleges gráfok irányítatlan pontösszefüggőség-növelésének bonyolultsága nyitott kérdés; az eggyel való növelés önmagában is sokat vizsgált terület. A harmadik részben bizonyított formula Frank és Jordán 1994-ből származó sejtése.
- *Megadjuk az első kombinatorikus polinomiális algoritmust irányított pontösszefüggőség-növelésre.* Erre a problémára Frank és Jordán 1995-ben adtak min-max formulát. Nyitott maradt azonban a kérdés: hogyan található meg egy optimális megoldás kombinatorikus algoritmus segítségével. Az értekezésben megadunk két, teljesen különböző kombinatorikus algoritmust. A második rész az összefüggőség eggyel való növelésének speciális esetét oldja meg algoritmikusan (Frank Andrással közös eredmény), a negyedik rész pedig az általános problémára ad algoritmust (ifj. Benczúr Andrással közös eredmény). Valójában még általánosabb problémát oldunk meg: új, algoritmikus bizonyítást adunk Frank és Jordán általános halmazpárfedési tételére is.
- *Megadjuk a (k, ℓ) -élösszefüggő gráfok egy konstruktív karakterizációját.* A hatodik részben bemutatott, Kovács Erika Renátával közös eredmény Frank 2003-as sejtését bizonyítja be. A tétel több korábbi karakterizáció közös általánosítását adja, és természetesen illeszkedik az eddig leemelési és irányítási tételek rendszerébe.
- *Részleges eredményeket adunk a partíciókorlátos irányítatlan lokális élösszefüggőség-növelési problémára.* Az ötödik részben irányítatlan élösszefüggőség-növeléssel kapcsolatban tárgyalunk néhány klasszikus eredményt egységes keretben, az élátbillentési technikát használva. A partíciókorlátos problémával kapcsolatban megfogalmazzuk és részben bebizonyítunk egy sejtést.

Az eredmények nagy része a [36], [74], [73] és [56] cikkekből származik. Kivételt képez az ötödik rész, amely nem publikált eredményeket tartalmaz.