

On some problems related to 2-level polytopes



Thèse n. 9025 2018
présenté le 22 Octobre 2018
à la Faculté des Sciences de Base
chaire d'optimisation discrète
programme doctoral en mathématiques
École Polytechnique Fédérale de Lausanne
pour l'obtention du grade de Docteur ès Sciences
par

Manuel Francesco Aprile

acceptée sur proposition du jury:

Prof Hess Bellwald Kathryn, président du jury
Prof Eisenbrand Friedrich, directeur de thèse
Prof Faenza Yuri, codirecteur de thèse
Prof Svensson Ola, rapporteur
Prof Fiorini Samuel, rapporteur
Prof Weltge Stefan, rapporteur

Lausanne, EPFL, 2018

"You must go from wish to wish.
What you don't wish for will always be beyond your reach."
— Michael Ende, *The Neverending Story*.

To my brother Stefano,
to whom I dedicated my Bachelor thesis too.
He is now six years old and he has not read it yet...
Maybe one day?

Acknowledgements

My PhD has been an incredible journey of discovery and personal growth. There were moments of frustration and doubt, when research just wouldn't work, even depression, but there was also excitement, enjoyment, and satisfaction for my progresses. Oh, there was hard work too. A whole lot of people contributed to making this experience so amazing.

First, I would like to thank my advisor Fritz Eisenbrand for his support, the trust he put in me, and for letting me do research in total freedom. He allowed me to go to quite some conferences and workshops, where I could expand my knowledge and professional network, as well as enjoying some traveling.

A big “grazie” goes to Yuri Faenza, my co-advisor: in a moment when I was not sure about the direction of my research, he gave a talk about 2-level polytopes, immediately catching my interest, and, a few days after, he was introducing me to extended formulations with an improvised crash course. Ironically, that day he explained to me Yannakakis' protocol, which became the core of the last chapter of my thesis. He invested many hours working with me, taught me so much, he even hosted me at his place in New York (thanks to his wife Fanny too!): it goes without saying that this thesis would not have been possible without him.

I am indebted to all researchers I met, from which I gathered interesting conversations, inspiration, precious pieces of advice. A complete list would be too long, but I would like to mention Michele Conforti, Marco Di Summa, Gwenaël Joret, Volker Kaibel, Monaldo Mastrolilli, Janos Pach. In particular I thank Kathryn Hess Bellwald, Ola Svensson, Stefan Weltge for being part of my PhD jury. Heartfelt thanks to Samuel Fiorini for inviting me to Bruxelles and for teaching me so much in so little time. Thanks to Tony Huynh and Marco Macchia, it was great to work and go for beers with you. I look forward to repeating both experiences. Thanks to Alfonso Cevallos, best office mate ever, for all the adventures, discussions, and for solving problems that I could not solve. Thanks to each member of DISOPT group, in particular to Igor Malinović and Christoph Hunkenschröder for their valuable friendship and all the fun activities we did together, and to the members of DCG group too. We had quite a good time! I am also very grateful to Jocelyne and Natacha for their kindness and precious administrative support. I also had the luck to make many friends in the math department and beyond, among which I would like to mention Riccardo (worst housemate!), Clelia, Francesco, Gennaro, Jessica, Arwa, Lydia, Thomas, Chiara, Mauro. They might not have directly contributed to my work, but definitely made my time off much more enjoyable.

Finally, the support and the love I had from Sicily was very important to me, especially during difficult moments. Thanks to all my Sicilian friends. Thanks to Carla, who made the last part of my PhD way better than the first in so many aspects: trying to explain my research to you

Acknowledgements

was so fun that I plan to do it again and again, you have been warned. I am much grateful to all my family and in particular to my parents, for being extremely supportive and for encouraging me to pursue my dreams with no limits but my ambition, even though this meant to go far from home. Grazie.

Lausanne, 17 September 2018

Abstract

In this thesis we investigate a number of problems related to 2-level polytopes, in particular regarding their combinatorial structure and extension complexity. 2-level polytopes have been introduced as a generalization of stable set polytopes of perfect graphs, and despite their apparently simple structure, are at the center of many open problems: these include connection with communication complexity and the separation between linear and semidefinite programming. The extension complexity of a polytope P is a measure of the complexity of representing P : it is the smallest size of an extended formulation of P , which in turn is a linear description of a polyhedron that projects down to P .

In the first chapter we introduce the main concepts that will be used through the thesis and we motivate our interest in 2-level polytopes.

In the second chapter we examine several classes of 2-level polytopes arising in combinatorial settings and we prove a relation between the number of vertices and facets of such polytopes, which is conjectured to hold for all 2-level polytopes. The proofs are obtained through an improved understanding of the combinatorial structure of such polytopes, which in some cases leads to results of independent interest.

In the third chapter we study the extension complexity of a restricted class of 2-level polytopes, the stable set polytopes of bipartite graphs, for which we obtain improved lower and upper bounds.

In the fourth chapter we study slack matrices of 2-level polytopes, important combinatorial objects related to extension complexity, defining operations on them and giving algorithms for the following recognition problem: given a matrix, determine whether it is a slack matrix of some special class of 2-level polytopes.

In the fifth chapter we address the problem of explicitly obtaining small size extended formulations whose existence is guaranteed by communication protocols. In particular we give an output-efficient algorithm to write down extended formulations for the stable set polytope of perfect graphs, making a well known result by Yannakakis constructive, and we extend this to all deterministic protocols.

We then conclude the thesis outlining the main open questions that stem from our work.

Keywords: Polytopes, polyhedral combinatorics, 2-level, extension complexity, vertices, facets, slack matrix.

Sommario

In questa tesi vengono trattati diversi problemi sui politopi "2-level", in particolare sulla loro struttura combinatoria e complessità di estensione. Tali politopi sono una generalizzazione di politopi che derivano dagli insiemi indipendenti nei grafi perfetti, e, nonostante la loro struttura apparentemente semplice, sono al centro di molti problemi aperti che spaziano dalla complessità computazionale alla programmazione semidefinita. La complessità di estensione di un politopo P è una misura della complessità nel rappresentare P : è la minima dimensione di una formulazione estesa di P , che a sua volta è una descrizione lineare di un poliedro di cui P è la proiezione.

Nel primo capitolo vengono introdotti i politopi 2-level e i concetti principali che verranno usati nella tesi, e vengono descritte le principali motivazioni dell'interesse verso questi politopi.

Nel secondo capitolo vengono esaminate diverse classi di politopi 2-level che appaiono in contesti combinatori, e viene provata una relazione tra il numero di faccette e di vertici di tali politopi. Congetturiamo che tale relazione valga per tutti i politopi 2-level. Le dimostrazioni vengono ottenute tramite una migliore comprensione della struttura combinatoria di tali politopi, che a volte porta a risultati interessanti a prescindere dalla congettura.

Nel terzo capitolo studiamo la complessità di estensione di una particolare classe di politopi 2-level, derivante dagli insiemi indipendenti dei grafi bipartiti, di cui miglioriamo il limite inferiore e superiore.

Nel quarto capitolo studiamo le matrici di slack dei politopi 2-level, importanti oggetti combinatori collegati alla complessità di estensione, definiamo operazioni su tali matrici e diamo algoritmi per il seguente problema: data una matrice, determinare se è una matrice di slack di una certa classe di politopi 2-level.

Nel quinto capitolo affrontiamo il problema di ottenere formulazioni estese compatte ed esplicite, quando l'esistenza di tali formulazioni è dimostrata tramite protocolli di comunicazione. In particolare diamo un algoritmo per ottenere una formulazione estesa del politopo degli insiemi indipendenti dei grafi perfetti, rendendo costruttivo un noto risultato di Yannakakis. Il risultato è abbastanza generale da essere applicabile a tutti i protocolli deterministici.

La tesi si conclude con una discussione delle principali direzioni di ricerca che scaturiscono dal nostro lavoro.

Parole chiave: politopi, 2-level, formulazioni estese, vertici, faccette, matrice di slack.

Contents

Acknowledgements	v
Abstract (English/Italiano)	vii
1 Introduction	1
1.1 Preliminaries	6
2 On vertices and facets of 2-level polytopes arising in combinatorial settings	7
2.1 Introduction	7
2.2 Basics	9
2.2.1 Hanner and Birkhoff polytopes	9
2.3 Graphical 2-Level Polytopes	10
2.3.1 Stable set polytopes of perfect graphs	11
2.3.2 Hansen polytopes	12
2.3.3 Min up/down polytopes	12
2.3.4 Polytopes coming from posets	14
2.3.5 Stable matching polytopes	16
2.4 2-Level Matroid Base Polytopes	19
2.4.1 2-level matroid polytopes and Conjecture 2.1	20
2.4.2 Facets of 2-sums	24
2.4.3 Linear Description of 2-Level Matroid Base Polytopes	27
2.5 Cut Polytope and Matroid Cycle Polytope	30
2.6 On possible generalizations of the conjecture	33
2.6.1 Forest polytope of $K_{2,n}$	33
2.6.2 Spanning tree polytope of the skeleton of the 4-dimensional cube	33
2.6.3 3-level min up/down polytopes	33
2.6.4 Polytopes of minimum PSD rank	34
2.6.5 Polytopes with structured linear relaxations	35
2.6.6 0/1 matrices generalizing slack matrices of 2-level polytopes	35
3 On the extension complexity of the stable set polytope of bipartite graphs	39
3.1 Introduction	39
3.2 Rectangle Covers and Fooling Sets	41
3.3 An Improved Upper Bound	42

Contents

3.4	An Improved Lower Bound	43
3.5	A small rectangle cover of the special entries	47
3.6	A connection with rectangle covers of the spanning tree polytope	51
3.6.1	The class covering problem	51
3.6.2	Special entries, centered rectangles and class covers	52
3.6.3	Spanning tree polytopes	53
4	Slack matrices of 2-level polytopes: recognition and decomposition	57
4.1	Introduction	57
4.1.1	Preliminaries	58
4.2	k -sums of slack-matrices	59
4.2.1	1-sums	59
4.2.2	2-sums and k -sums	61
4.3	Slack matrices of 2-level stable set polytopes	66
4.4	Recognition algorithms	68
4.4.1	Recognizing 1-sums via submodular function minimization	69
4.4.2	Extension to k -sums	73
4.4.3	Numerical experiments	74
4.4.4	Matroid polytopes	75
4.5	Matroid polytopes: an alternative approach	80
4.5.1	Phase 1: finding the circuits of M	81
4.5.2	Phase 2: reconstructing $B(M)$ and its slack matrix	88
5	Extended formulations in output-efficient time from communication protocols	91
5.1	Introduction	91
5.2	Preliminaries	92
5.2.1	Deterministic and non-deterministic protocols	92
5.2.2	Extended formulations for a pair of polytopes	93
5.2.3	Protocols and extended formulations	94
5.2.4	The stable set polytope	95
5.3	A general approach	96
5.3.1	Application to $(STAB(G), QSTAB(G))$	98
5.4	Direct derivations	100
5.4.1	Complement graphs	100
5.4.2	Alternative formulation for $STAB(G)$, G perfect	101
5.4.3	Claw-free perfect graphs and generalizations	103
5.4.4	Comparability graphs	105
6	Conclusion	107
A	An appendix	111
A.1	The polytopes from Proposition 2.11, Chapter 2	111
A.2	The polytope from Section 2.6.5	113

Bibliography	121
Curriculum Vitae	123

1 Introduction

A classical, powerful approach in discrete optimization is to represent the feasible solutions of a problem as vertices of a polytope and to use linear programming to find the optimal vertex. Hence, a solid mathematical understanding of polytopes associated to combinatorial problems is a fundamental goal of the modern theory of optimization. In this thesis we study a number of problems concerning a particular class of polytopes, called 2-level. Such polytopes have an apparently simple structure and appear in several different contexts; yet, our understanding of them is relatively poor. This makes them fascinating objects, especially from the point of view of optimization.

Definition 1.1. A polytope $P \subset \mathbb{R}^d$ is called *2-level* if, for any supporting hyperplane H defining a facet F , there is a hyperplane parallel to H that contains all the vertices of P that are not in F .




Figure 1.1 – The first three polytopes (the simplex, the cross-polytope and the cube) are 2-level. The fourth one is not 2-level, because of the highlighted facet.

2-level polytopes naturally arise in many areas of mathematics, and they were defined independently in at least two different contexts:

- Sum of squares and polynomial ideals: in [45] the Theta body of the real variety of an ideal is introduced as a relaxation based on sum of squares, and 2-level polytopes are defined as those polytopes for which this relaxation is exact.
- Statistics: in [93] a polytope is called *compressed* if all its pulling triangulations are

unimodular with respect to the lattice generated by the vertices, and in [94] this property is shown to be equivalent to being 2-level.

The property of 2-levelness, although quite strong, is satisfied by several classes of polytopes: Birkhoff [101], Hanner [54], order polytopes [92], spanning tree polytopes of series-parallel graphs [48], stable matching polytopes [52], and most importantly stable set polytopes of perfect graphs [45], which are discussed below. It is not a coincidence that the aforementioned polytopes have 0/1 vertices: in [45] it is shown that each 2-level polytope is affinely isomorphic to a 0/1 polytope (i.e. a polytope whose vertices have 0/1 coordinates). This implies that there is a finite number of (equivalence classes of) 2-level polytopes of a given dimension d , in particular at most 2^{2^d} . However, 2-level polytopes seem to form a very restricted and in some sense well-behaved subclass of 0/1 polytopes. For instance, it is not hard to see that every face of a 2-level polytope is again 2-level. Using this and other structural results, in [10] an algorithm is given for enumerating 2-level polytopes, and a complete enumeration is done up to dimension 6 (this was extended to dimension 7 and 8 in subsequent versions [11, 76]). In the paper it is argued that for general 0/1 polytopes such a task is not practically feasible, as with the current computational power one cannot even store all the equivalence classes already for dimension 6. This enumeration showed that already in low dimensions there are many 2-level polytopes that do not have an immediate combinatorial interpretation and are outside the classes described above, suggesting that an understanding of all 2-level polytopes that goes beyond the special cases is desirable. Based on their experimental evidence, the authors of [11] conjectured that the number of 2-level polytopes of dimension d is at most $2^{\text{poly}(d)}$. This was recently proved in [34], where an upper bound of $2^{O(d^2 \log d)}$ has been given, together with an almost-matching lower bound of $2^{\Omega(d^2)}$.

While 2-level polytopes seem to be a "small" class, there are many open questions about them. In particular we now define a parameter that is a current theme of this thesis for its relation to 2-level polytopes, namely extension complexity. In the context of optimization, it is crucial to have compact representations of our polytopes of interest. Obtaining our polytope P as a projection of a higher dimensional polyhedron Q can drastically reduce the size of the representation and make the problem of optimizing over P efficiently solvable.

Definition 1.2. Let $P \in \mathbb{R}^n$ be a polytope. A polyhedron $Q \subseteq \mathbb{R}^p$ is an *extension* of P if there exists an affine map $\pi : \mathbb{R}^p \rightarrow \mathbb{R}^d$ with $\pi(Q) = P$. An *extended formulation* of P is a linear description of an extension of a P , and the *extension complexity* of P (denoted by $xc(P)$) is the smallest number of facets of any extension of P (equivalently, the smallest number of inequalities in any extended formulation of P).

One can also consider *semidefinite extended formulations*: in this case Q , instead of a polyhedron, is an affine slice of the semidefinite cone, again with the requirement that the projection on the original space is P , and the *semidefinite extension complexity* of P is the minimum dimension of the semidefinite cone in any such Q . Although semidefinite programming can be seen as a generalization of linear programming and it is solvable in polynomial time up

to arbitrary precision by interior point methods, in practice solving linear programs is more efficient and preferable (see for instance [69]). Hence there is interest in finding small (linear) extended formulations even when semidefinite formulations are already given, which is the case for 2-level polytopes.

One of the reasons of interest in the extension complexity of 2-level polytopes comes from the fact that they were introduced as a generalization of stable set polytopes of perfect graphs. The class of perfect graphs has received much attention in the literature since the 1960s, when Berge introduced them and formulated a conjecture on them [8]. After more than forty years of partial results this conjecture was proved by Chudnovsky, Robertson, Seymour and Thomas, under the name of Strong Perfect Graph Theorem [15]. Perfect graphs have quite special properties in terms of their cliques and stable sets (also called independent sets). In particular they can be characterized in terms of their stable set polytope, which has the following simple (yet exponential in size) description:

$$\text{STAB}(G) = \left\{ x \in \mathbb{R}_+^d : \sum_{v \in C} x_v \leq 1 \text{ for all maximal cliques } C \text{ of } G \right\},$$

where G is a perfect graph on d vertices. From this description (due to Chvátal, [16]) it is easy to see that this polytope is 2-level: indeed, for any vertex x and for any clique C the quantity $\sum_{v \in C} x_v$ can be either 0 or 1, giving two parallel hyperplanes that contain all the vertices for every facet defining direction. Moreover it can be shown that, for a graph G , $\text{STAB}(G)$ is 2-level if and only if G is perfect [45].

In [73] Lovász introduced the *Theta body* of a graph G as a convex body that approximates $\text{STAB}(G)$. If G has n vertices, its Theta body can be expressed by a semidefinite program of size $n + 1$, and if G is perfect, its Theta body is an exact semidefinite formulation of $\text{STAB}(G)$, hence we can efficiently find a maximum weight stable set in G using semidefinite programming. To find a purely combinatorial algorithm for this problem is the most important open question on perfect graphs. Thirty years later, in [45], the concept of Theta body was extended to define a hierarchy of semidefinite relaxations (the k -th Theta body, with k a positive integer) to approximate the convex hull of any set of points. As already mentioned, 2-level polytopes can be characterized as those polytopes for which the first level of this hierarchy is exact, i.e. in particular they have small semidefinite extension complexity. This implies that, in principle, one can optimize over these polytopes in polynomial time using semidefinite programming, generalizing Lovász's result to all 2-level polytopes. The question left open is whether we can achieve the same using linear programming only, i.e. what is the extension complexity of 2-level polytopes. Whereas in general the best upper bound known is superpolynomial in the dimension ([75]), for stable set polytopes of perfect graphs a quasipolynomial bound was given by Yannakakis (see Theorem 5.4, or [100]). Whether this is tight or can be improved to a polynomial bound is a prominent open question, as it is open whether the bound can be extended to all 2-level polytopes. Moreover, it is not known (see [35]) whether there is a polytope with exponential (or superpolynomial) extension complexity and polynomial

semidefinite extension complexity. This is a fundamental question: how much more powerful is semidefinite programming than linear programming? This has been answered in some cases: for instance, semidefinite programming gives a better approximation ratio for the Max-Cut problem than any known linear program (see [goemans1995improved]) and this extends to more general settings (see the superiority of Lasserre hierarchy over Sherali-Adams [71]). 2-level polytopes are perfect candidates to answer this question, as they have compact semidefinite complexity, but might have superpolynomial extension complexity.

In [100], Yannakakis showed that extension complexity of a polytope is captured by its slack matrix, defined as follows:

Definition 1.3. Given a polytope P described as $P = \text{conv}(v_1, \dots, v_n) = \{x \in \mathbb{R}^d : Ax \leq b\}$, where A has m rows, the *slack matrix* $S(P)$ is a non-negative $m \times n$ matrix with $S(P)_{i,j} = b_i - a_i^\top v_j$, i.e., the (i, j) -th entry is the slack of point j with respect to the i -th inequality.

Notice that the slack matrix of a given polytope is not uniquely determined as it depends on the vertical (also called inner) and horizontal (outer) representation that we choose. However in most cases the properties of interest of such matrices do not depend on the representation, hence it makes sense to refer to *the* slack matrix of a polytope. Slack matrices have interesting geometrical properties and the problem of determining whether a given matrix is a slack matrix is equivalent to the Polyhedral Verification problem, whose computational complexity is unknown [43]. Notice that, as a direct consequence of Definition 1.1, 2-level polytopes can be characterized as those polytopes having the “simplest” slack matrices.

Observation 1.4. *Let P be a polytope, then P is 2-level if and only if it admits a slack matrix with 0/1 entries only.*

Yannakakis’ work relates the extension complexity of a polytope to nonnegative factorizations of its slack matrix. The *nonnegative rank* of M is the smallest intermediate dimension in a nonnegative factorization of M , i.e. the smallest r such that there exist $T \in \mathbb{R}_{\geq 0}^{m \times r}$, $U \in \mathbb{R}_{\geq 0}^{r \times n}$ with $M = TU$.

Theorem 1.5. [100] *Given a polytope P of dimension at least 1 and its slack matrix S , the extension complexity of P is equal to the nonnegative rank of S .*

Yannakakis’ Theorem allows to study extension complexity with tools from linear algebra and combinatorics. This, as we will describe in Chapter 5, implies a beautiful connection between extension complexity and communication complexity: the latter field aims at understanding the amount of information that needs to be exchanged between two parties in order to compute a matrix given as input. This matrix usually expresses a predicate and is 0/1, in particular the *log-rank conjecture*, a fundamental open problem in the field, is concerned with the deterministic communication complexity of such matrices (we refer to [75] for more details). In light of Observation 1.4, 2-level polytopes are directly related to the log-rank conjecture. If true,

the conjecture would imply that the extension complexity of a d -dimensional 2-level polytope is at most $2^{\text{polylog}(d)}$, hence quasipolynomial. The best bound that is currently known is $2^{O(\sqrt{d})}$, and it is implied by the work of Lovett [75] on the log-rank conjecture. This means that, while no 2-level polytope can have exponential extension complexity, it might be possible to prove a lower bound of the kind $2^{\Omega(n^c)}$ for some $c \leq 1/2$. A 2-level polytope exhibiting such a bound would refute the log-rank conjecture.

For the reasons cited above, finding upper and lower bounds on the extension complexity of 2-level polytopes is a problem of prominent interest, arguably one of the biggest problems that are left open in the field. The major obstacle is that we lack a complete understanding of 2-level polytopes, of their combinatorial properties and geometric structure. In this thesis we investigate a number of problems related to 2-level polytopes, their slack matrices and their extension complexity, with a two-fold aim: to give contributions to the open questions cited so far; to expand our current knowledge on 2-level polytopes and propose new perspectives and tools for their study. Although the problems that we examine usually focus on special classes of 2-level polytopes, arising from combinatorial objects like graphs and matroids, we hope that some of techniques used can be extended to more general settings. On the way to our proofs, we also give contributions whose interest goes beyond 2-level polytopes: most notably we obtain results on the number of cliques and stable sets in a graph (Section 2.3), on the structure of the matroid base polytope (Sections 2.4.2, 4.5), on a combinatorial problem related to the spanning tree polytope (Section 3.6) and on extended formulations of general polytopes (Chapter 5).

The rest of the thesis is organized as follows:

- In Chapter 2, we examine a conjecture posed in [10] on the number of vertices and facets of 2-level polytopes and prove that it holds for many classes of 2-level polytopes coming from combinatorial settings. In doing so, we obtain a number of results which shed light on the structure of some 2-level polytopes, most notably stable marriage polytopes and matroid polytopes. The content of this chapter is joint work with Alfonso Cevallos and Yuri Faenza, and it has appeared, with some modifications, in [2] and [1].
- In Chapter 3, we study the extension complexity of stable set polytopes of bipartite graphs. In particular we derive the first non-trivial lower bound on the extension complexity of such polytopes, which is also the first lower bound for general 2-level polytopes. We show that our lower bound cannot be improved by using our technique, and in doing so we outline a connection with the extension complexity of the spanning tree polytope. The content of this chapter is joint work with Yuri Faenza, Samuel Fiorini, Tony Huynh, Marco Macchia and appears in [3], except for Section 3.6, which is joint work with Jana Cslovjcek.
- In Chapter 4, we study 0/1 slack matrices, i.e. slack matrices of 2-level polytopes, and the algorithmic problem of recognizing such matrices efficiently. In particular we introduce some operations on slack matrices that preserve 2-levelness and allow, thanks to a

decomposition approach, to recognize slack matrices of 2-level matroid polytopes. The content of this chapter is joint work with Michele Conforti, Yuri Faenza, Samuel Fiorini, Tony Huynh, Marco Macchia.

- In Chapter 5 we examine the algorithmic problem of obtaining extended formulations in output-efficient time, when the existence of such formulation is guaranteed by a communication protocol. In particular we focus on the stable set polytope of perfect graphs and we turn Yannakakis' quasipolynomial bound, mentioned above, into a quasipolynomial time algorithm. We also extend this to the more general setting of deterministic protocols, going beyond 2-level polytopes. This is joint work with Yuri Faenza and Mihalis Yannakakis.
- In Chapter 6 we conclude by describing further research directions and open questions left by our work.

1.1 Preliminaries

We let \mathbb{R}_+ be the set of nonnegative real numbers. For a set S and an element e , we denote by $A + e$ and $A - e$ the sets $A \cup \{e\}$ and $A \setminus \{e\}$, respectively. For a point $x \in \mathbb{R}^I$, where I is an index set, and a subset $J \subseteq I$, we let $x(J) = \sum_{i \in J} x_i$.

For a polytope $P \in \mathbb{R}^d$, we denote by $f_k(P)$ the number of k -dimensional faces of P . The polar of P is the polyhedron $P^\Delta = \{y \in \mathbb{R}^d : y \cdot x \leq 1 \forall x \in P\}$. It is well known¹ that, if $P \subseteq \mathbb{R}^d$ is a d -dimensional polytope with the origin in its interior, then so is P^Δ , and one can define a one-to-one mapping between vertices (resp. facets) of P and facets (resp. vertices) of P^Δ . The d -dimensional *cube* is $[-1, 1]^d$, and the d -dimensional *cross-polytope* is its polar.

One of the most common operation with polytopes is the *Cartesian product*. Given two polytopes $P_1 \subseteq \mathbb{R}^{d_1}$, $P_2 \subseteq \mathbb{R}^{d_2}$, their Cartesian product is $P_1 \times P_2 = \{(x, y) \in \mathbb{R}^{d_1+d_2} : x \in P_1, y \in P_2\}$.

¹It immediately follows from e.g. [101, Theorem 2.11].

2 On vertices and facets of 2-level polytopes arising in combinatorial settings

2.1 Introduction

In this chapter we present a polyhedral study of 2-level polytopes arising from combinatorial settings. In particular, the number of vertices and facets of such polytopes is studied. Each d -dimensional 2-level polytope is affinely isomorphic to a 0/1 polytope [45], hence it has at most 2^d vertices. Interestingly, the authors of [45] also showed that a d -dimensional 2-level polytope also has at most 2^d facets. This makes 2-level polytopes quite different from “random” 0/1 polytopes, that have $(d/\log d)^{\Theta(d)}$ facets [7]. Experimental results from [10, 76] suggest that this separation could be even stronger: up to $d = 8$, the product of the number of facets $f_{d-1}(P)$ and the number of vertices $f_0(P)$ of a d -dimensional 2-level polytope P does not exceed $d2^{d+1}$. In [10], it is asked whether this always holds, and in their journal version the question is turned into a conjecture.

Conjecture 2.1 (Vertex/facet trade-off). Let P be a d -dimensional 2-level polytope. Then

$$f_0(P)f_{d-1}(P) \leq d2^{d+1}.$$

Moreover, equality is achieved if and only if P is affinely isomorphic to the cross-polytope or the cube.

It is immediate to check that the cube and the cross-polytope (its polar) indeed verify $f_0(P)f_{d-1}(P) = d2^{d+1}$. Conjecture 2.1 has an interesting interpretation as an upper bound on the “size” of slack matrices of 2-level polytopes, since $f_0(P)$ (resp. $f_{d-1}(P)$) is the number of columns (resp. rows) of the (smallest) slack matrix of P . Many fundamental results on linear extensions of polytopes are based on properties of their slack matrices. We believe that advancements on Conjecture 2.1 may lead to precious insights on the structure of (the slack matrices of) 2-level polytopes, similarly to how progresses on e.g. the outstanding Hirsch [88] and 3^d conjectures for centrally symmetric polytopes [62] shed some light on our general understanding of polytopes.

Contribution and organization.

The main results of this chapter are the following:

- We give considerable evidence supporting Conjecture 2.1 by proving it for several classes of 2-level polytopes arising in combinatorial settings. These include polytopes coming from graphs (stable set, Hansen, and stable matching polytopes), from posets (Order, Chain and double order polytopes) and from matroids (base matroid and cycle polytopes) and Birkhoff, Hanner and min up-down polytopes. We refer to the following sections for relevant definitions and references.
- We establish new properties of many classes of 2-level polytopes, of their underlying combinatorial objects, and of their inter-class connections. These results include: a trade-off formula for the number of stable sets and cliques in a graph; a description of the stable matching polytope as an affine projection of the order polytope of the associated rotation poset; a non-redundant characterization of facet-defining inequalities for base polytopes of matroids under the 2-sum operation; and a compact linear description of 2-level base polytopes of matroids in terms of cuts of some trees associated to those matroids (notably, our description has linear size in the dimension and can be written down explicitly in polynomial time). These results simplify the algorithmic treatment of some of these polytopes, as well as provide a deeper combinatorial understanding of them. At a more philosophical level, these examples suggest that being 2-level is a very attractive feature for a (combinatorial) polytope, since it seems to imply a well-behaved underlying structure.
- We moreover show examples of 0/1 polytopes with a simple structure (including spanning tree and forest polytopes) that are not 2-level and *do not* satisfy Conjecture 2.1. This suggests that, even though there are clearly polytopes that are not 2-level and satisfy Conjecture 2.1, 2-levelness seem to be the “correct” hypothesis to prove a general positive result. We also investigate extensions of the conjecture in terms of matrices and systems of linear inequalities.

We introduce some basic definitions and techniques in Section 2.2: those are enough to show that Conjecture 2.1 holds for Birkhoff and Hanner polytopes. In Section 2.3, we first prove an upper bound on the product of the number of stable sets and cliques of a graph (see Theorem 2.5). We then prove Conjecture 2.1 for stable set polytopes of perfect graphs, Hansen polytopes, min up-down polytopes, order, double order and chain polytopes of posets, and stable matching polytopes, by reducing these results to statements on stable sets and cliques of associated graphs, which are also proved in Section 2.3. Hence, we call all those *graphical* 2-level polytopes. Of particular interest is our observation that stable matching polytopes are affine equivalent to order polytopes (see Theorem 2.18). In Section 2.4, we study 2-level matroid base polytopes, and prove that Conjecture 2.1 for this class (see Theorem 2.26). The section also includes results on base polytopes of general matroids (see Theorem 2.30, Corollary 2.31), which we believe of independent interest. Using this results, we derive a compact description of 2-level base polytopes of matroids (see Theorem 2.35). In Section 2.5, we prove the conjecture for the cycle polytopes of certain binary matroids, which generalizes

all cut polytopes that are 2-level. In Section 2.6, we investigate possible extensions of the conjecture.

2.2 Basics

Let $P \in \mathbb{R}^d$ a polytope, and P^Δ its polar, as defined in Section 1.1. Since $f_0(P) = f_{d-1}(P^\Delta)$, and $f_0(P^\Delta) = f_{d-1}(P)$, a polytope and its polar will simultaneously satisfy or not satisfy Conjecture 2.1. Recall that a 0/1 polytope is the convex hull of a subset of the vertices of $\{0, 1\}^d$. The following facts will be used many times:

Lemma 2.2. [45] *Let P be a 2-level polytope of dimension d . Then*

1. $f_0(P), f_{d-1}(P) \leq 2^d$.
2. *Any face of P is again a 2-level polytope.*

As a preliminary observation we show that the operation of Cartesian product preserves 2-levelness and the bound of Conjecture 2.1.

Lemma 2.3. *Two polytopes P_1, P_2 are 2-level if and only if their Cartesian product $P_1 \times P_2$ is 2-level. Moreover, if two 2-level polytopes P_1 and P_2 satisfy Conjecture 2.1, then so does $P_1 \times P_2$.*

Proof. The first part follows immediately from the fact that $P_1 = \{x : A^{(1)}x \leq b^{(1)}\}$, $P_2 = \{y : A^{(2)}y \leq b^{(2)}\}$, then $P_1 \times P_2 = \{(x, y) : A^{(1)}x \leq b^{(1)}; A^{(2)}y \leq b^{(2)}\}$, and that the vertices of $P_1 \times P_2$ are exactly the points (x, y) such that x is a vertex of P_1 and y a vertex of P_2 .

For the second part, let $P = P_1 \times P_2$, $d_1 = d(P_1)$, $d_2 = d(P_2)$. Then it is well known that $d(P) = d_1 + d_2$, $f_0(P) = f_0(P_1)f_0(P_2)$, and $f_{d-1}(P) = f_{d_1-1}(P_1) + f_{d_2-1}(P_2)$. We conclude

$$\begin{aligned} f_0(P)f_{d-1}(P) &= f_0(P_1)f_{d_1-1}(P_1)f_0(P_2) + f_0(P_2)f_{d_2-1}(P_2)f_0(P_1) \\ &\leq d_1 2^{d_1+d_2+1} + d_2 2^{d_1+d_2+1} \\ &= d(P) 2^{d(P)+1}, \end{aligned}$$

where the inequality follows by induction and from Lemma 2.2. Suppose now that P satisfies the bound with equality. Then, for $i = 1, 2$, P_i also satisfies the bound with equality and $f_0(P_i) = 2^{d(P_i)}$, which means that P_i is a d_i -dimensional cube. Then P is a d -dimensional cube. \square

2.2.1 Hanner and Birkhoff polytopes

We start off with two easy examples. *Hanner* polytopes [53] are defined as the smallest family that contains the $[-1, 1]$ segment of dimension 1, and is closed under taking polars and

Cartesian products. That they verify the conjecture immediately follows from Lemma 2.3 and from the discussion on polars earlier in Section 2.2. The *Birkhoff* polytope $B_n \subset \mathbb{R}^{n^2}$ is the convex hull of all $n \times n$ permutation matrices (see e.g. [101]). For $n = 2$, the polytope B_2 is affinely isomorphic to the Hanner polytope of dimension 1. For $n \geq 3$, B_n is known [101] to have exactly $n!$ vertices, n^2 facets, dimension $(n - 1)^2$, and is 2-level. We conclude the following.

Lemma 2.4. *Hanner and Birkhoff polytopes satisfy Conjecture 2.1.*

2.3 Graphical 2-Level Polytopes

We present a general result on the number of cliques and stable sets of a graph. Proofs of all theorems from the current section will be based on it.

Theorem 2.5 (Stable set/clique trade-off). *Let $G = (V, E)$ be a graph on n vertices, \mathcal{C} its family of non-empty cliques, and \mathcal{S} its family of non-empty stable sets. Then*

$$|\mathcal{C}||\mathcal{S}| \leq n(2^n - 1).$$

Moreover, equality is achieved if and only if G or its complement is a clique.

Proof. Consider the function $f : \mathcal{C} \times \mathcal{S} \rightarrow 2^V$, where $f(C, S) = C \cup S$. For a set $W \subset V$, we bound the size of its pre-image $f^{-1}(W)$. If W is a singleton, the only pair in its pre-image is (W, W) . For $|W| \geq 2$, we claim that $|f^{-1}(W)| \leq 2|W|$.

There are at most $|W|$ intersecting pairs (C, S) in $f^{-1}(W)$. This is because the intersection must be a single element, $C \cap S = \{v\}$, and once it is fixed every element adjacent to v must be in C , and every other element must be in S .

There are also at most $|W|$ disjoint pairs in $f^{-1}(W)$, as we prove now. Fix one such disjoint pair (C, S) , and notice that both C and S are non-empty proper subsets of W . All other disjoint pairs (C', S') are of the form $C' = C \setminus A \cup B$ and $S' = S \setminus B \cup A$, where $A \subseteq C$, $B \subseteq S$, and $|A|, |B| \leq 1$. Let X (resp. Y) denote the set formed by the vertices of C (resp. S) that are anticomplete to S (resp. complete to C). Clearly, either X or Y is empty. We settle the case $Y = \emptyset$, the other being similar. In this case $\emptyset \neq A \subseteq X$, so $X \neq \emptyset$. If $X = \{v\}$, then $A = \{v\}$ and we have $|S| + 1$ choices for B , with $B = \emptyset$ possible only if $|C| \geq 2$, because we cannot have $C' = \emptyset$. This gives at most $1 + |S| + |C| - 1 \leq |W|$ disjoint pairs (C', S') in $f^{-1}(W)$. Otherwise, $|X| \geq 2$ forces $B = \emptyset$, and the number of such pairs is at most $1 + |X| \leq 1 + |C| \leq |W|$.

We conclude that $|f^{-1}(W)| \leq 2|W|$, or one less if W is a singleton. Thus

$$|\mathcal{C} \times \mathcal{S}| \leq \sum_{k=0}^n 2k \binom{n}{k} - n = n2^n - n,$$

where the (known) fact $\sum_{k=0}^n 2k \binom{n}{k} = n2^n$ holds since

$$n2^n = \sum_{k=0}^n (k + (n-k)) \binom{n}{k} = \sum_{k=0}^n k \binom{n}{k} + (n-k) \binom{n}{n-k} = 2 \sum_{k=0}^n k \binom{n}{k}.$$

The bound is clearly tight for $G = K_n$ and $G = \overline{K_n}$. For any other graph, there is a subset W of 3 vertices that induces 1 or 2 edges. In both cases, $|f^{-1}(W)| = 5 < 2|W|$, hence the bound is loose. \square

Corollary 2.6. *Let G , \mathcal{C} and \mathcal{S} be as in Theorem 2.5, and $\mathcal{C}' = \mathcal{C} \cup \{\emptyset\}$ and $\mathcal{S}' = \mathcal{S} \cup \{\emptyset\}$ be the families of (possibly empty) cliques and stable sets of G , respectively. Then*

$$|\mathcal{C}'||\mathcal{S}'| \leq (n+1)2^n,$$

and equality is achieved if and only if G or its complement is a clique.

Proof. We apply the previous inequality to obtain

$$\begin{aligned} |\mathcal{C}'||\mathcal{S}'| &= (|\mathcal{C}| + 1)(|\mathcal{S}| + 1) = |\mathcal{C}||\mathcal{S}| + (|\mathcal{C}| + |\mathcal{S}'|) \\ &\leq n(2^n - 1) + (|\mathcal{C} \cup \mathcal{S}'| + |\mathcal{C} \cap \mathcal{S}'|) \\ &\leq n(2^n - 1) + (2^n + n) = (n+1)2^n. \end{aligned}$$

Clearly the inequality is tight whenever G or its complement is a clique, and from Theorem 2.5, we know that it is loose otherwise. \square

2.3.1 Stable set polytopes of perfect graphs

For a graph $G = (V, E)$, its stable set polytope $\text{STAB}(G)$ is the convex hull of the incidence vectors of the stable sets of G . We recall that $\text{STAB}(G)$ is 2-level if and only if G is a perfect graph [45], or equivalently [16] if and only if

$$\text{STAB}(G) = \{x \in \mathbb{R}_+^V : x(C) \leq 1 \text{ for all maximal cliques } C \text{ of } G\}.$$

Proposition 2.7. *Stable set polytopes of perfect graphs satisfy Conjecture 2.1.*

Proof. For a perfect graph $G = (V, E)$ on d vertices, the polytope $\text{STAB}(G)$ is d -dimensional. If we define \mathcal{C} , \mathcal{C}' and \mathcal{S}' as in Corollary 2.6, then the number of vertices in $\text{STAB}(G)$ is at most $|\mathcal{S}'|$. There are at most d non-negativity constraints, and at most $|\mathcal{C}| = |\mathcal{C}'| - 1$ clique constraints, so the number of facets in $\text{STAB}(G)$ is at most $|\mathcal{C}'| + d - 1$. Hence

$$\begin{aligned} f_0(\text{STAB}(G)) f_{d-1}(\text{STAB}(G)) &\leq (|\mathcal{C}'| + d - 1) |\mathcal{S}'| \\ &= |\mathcal{C}'||\mathcal{S}'| + (d - 1) |\mathcal{S}'| \\ &\leq (d + 1)2^d + (d - 1)2^d = d2^{d+1}, \end{aligned}$$

where we used Corollary 2.6 and the trivial inequality $|\mathcal{S}'| \leq 2^d$. We see that the conjectured inequality is satisfied, and is tight only in the trivial cases $d = 1$ or $|\mathcal{S}'| = 2^d$. In the latter case, G has no edges and $\text{STAB}(G)$ is affinely isomorphic to the cube. \square

2.3.2 Hansen polytopes

Given a $(d - 1)$ -dimensional polytope P , the *twisted prism* of P is the d -dimensional polytope defined as the convex hull of $\{(x, 1) : x \in P\}$ and $\{(-x, -1) : x \in P\}$. For a perfect graph G with $d - 1$ vertices, its *Hansen polytope* [54], $\text{Hans}(G)$, is defined as the twisted prism of $\text{STAB}(G)$. Hansen polytopes are 2-level and centrally symmetric, see e.g. [10].

Proposition 2.8. *Hansen polytopes satisfy Conjecture 2.1.*

Proof. Let $G = (V, E)$ be a perfect graph on $d - 1$ vertices, and let \mathcal{C}' and \mathcal{S}' be as in Corollary 2.6. Then $\text{Hans}(G)$ has $2|\mathcal{S}'|$ vertices (from the definition), and $2|\mathcal{C}'|$ facets (see e.g. [54]). Using again Corollary 2.6, we get

$$f_0(\text{Hans}(G)) f_{d-1}(\text{Hans}(G)) = 4|\mathcal{S}'||\mathcal{C}'| \leq 4d2^{d-1} = d2^{d+1}.$$

The inequality is tight only if G is either a clique or an anti-clique. The Hansen polytopes of these graphs are affinely equivalent to the cross-polytope and cube, respectively. \square

2.3.3 Min up/down polytopes

Fix two integers $0 < l < d$. For a 0/1 vector $x \in \{0, 1\}^d$ and index $1 \leq i \leq d - 1$, we call i a *switch index* of x if $x_i \neq x_{i+1}$. The vector x satisfies the *min up/down constraint* (with parameter l) if for any two switch indices $i < j$ of x , we have $j - i \geq l$. In other words, when x is seen as a bit-string then it consists of blocks of 0's and 1's each of length at least l (except possibly for the first and last blocks). The *min up/down polytope* $P_d(l)$ is defined as the convex hull of all 0/1 vectors in \mathbb{R}^d satisfying the min up/down constraint with parameter l . Those polytopes have been introduced in [72] in the context of discrete planning problems with machines that have a physical constraint on the frequency of switches between the operating and not operating states.¹ In [72, Theorem 4], the following characterization of the facet-defining inequalities of $P_d(l)$ is given.

Lemma 2.9. *Let $I \subset [d]$ be an index subset with elements $1 \leq i_1 < i_2 < \dots < i_k \leq d$, such that a) $k = |I|$ is odd and b) $i_k - i_1 \leq l$. Then, the two inequalities $0 \leq \sum_{j=1}^k (-1)^{j-1} x_{i_j} \leq 1$ are facet-defining for $P_d(l)$. Moreover, each facet-defining inequality in $P_d(l)$ can be obtained in this way.*

¹The more general definition given in [72] considers two parameters ℓ_1 and ℓ_2 , which respectively restrict the minimum lengths of the blocks of 0's and 1's in valid vertices. The resulting polytope is 2-level precisely when $\ell_1 = \ell_2$, thus in this section we restrict our attention to this case. General (non-2-level) min up/down polytopes do not satisfy Conjecture 2.1; see Example 2.6.3.

It is clear from this result that $P_d(l)$ is a 2-level polytope. Indeed, if all vertices of a polytope have 0/1 coordinates and all facet-defining inequalities can be written as $0 \leq c^T x \leq 1$ for integral vectors c , then the polytope is 2-level.

Proposition 2.10. *2-level min up/down polytopes satisfy Conjecture 2.1.*

Proof. Consider the 2-level min up/down polytope $P_d(l)$, for integers $0 < l < d$. $P_d(l)$ is full dimensional, hence it has dimension d . Define the graph $G([d-1], E)$, where $\{i, j\} \in E$ whenever $|j - i| \leq l - 1$, and let \mathcal{C}' and \mathcal{S}' be as in Corollary 2.6. We delay for a moment the proof of the following facts: a) $f_0(P_d(l)) = 2|\mathcal{S}'|$; and b) $f_{d-1}(P_d(l)) = 2|\mathcal{C}'|$. We obtain:

$$f_0(P_d(l))f_{d-1}(P_d(l)) = 4|\mathcal{S}'||\mathcal{C}'|.$$

This is the same inequality that appears in the proof of Proposition 2.8, hence in a similar fashion we conclude that the conjectured inequality is satisfied, and it is tight only if G is either a clique or an anti-clique. These cases correspond to $l = d - 1$ and $l = 1$, respectively, and it can be checked that $P_d(l)$ is then affinely equivalent to the cross-polytope or the cube.

Proof of fact a). For a vector $x \in \{0, 1\}^d$, let $I_x \subseteq [d-1]$ be its set of switch indices. Then x is (a vertex) in $P_d(l)$ iff I_x is a stable set in G . Moreover, if two vertices $x, y \in P_d(l)$ have exactly the same switch indices, then either $x = y$ or $x + y = \mathbf{1}$ (the all-ones vector). Hence, there is a mapping from the set of vertices of $P_d(l)$ to \mathcal{S}' , where each pre-image contains 2 elements. This proves the claim.

Proof of fact b). Let $\mathcal{S} \subseteq 2^{[d]}$ be the collection of all index sets $I \subseteq [d]$ satisfying the properties of Lemma 2.9. The lemma asserts that $f_{d-1}(P_d(l)) = 2|\mathcal{S}|$. To complete the proof, we present a bijection from \mathcal{S} to \mathcal{C}' . For $I \subset [d]$ in \mathcal{S} , let i be the lowest index in I , let $j = \min\{i + l, d\}$, and define $I' = I \setminus \{j\}$. I' is a clique in G . We conclude the proof by showing that the mapping can be inverted, hence it is bijective. Recall that G has nodes indexed from 1 to $d - 1$. For $I' \in \mathcal{C}'$, if $|I'|$ is odd, let $I = I'$; if $I' = \emptyset$, let $I = \{d\}$; otherwise, let i be the lowest index in I and $j = \min\{i + l, d\}$, and define $I = I' \cup \{j\}$. Clearly, in all cases $I \in \mathcal{S}$, and the preimages of two even cliques or two odd cliques are distinct. Now pick an even clique I' . If $I' = \emptyset$, then $I = \{d\}$ is not the preimage of an odd clique. If $I' \neq \emptyset$ and $i + l < d$, then I is not a clique of G , hence, in particular, it cannot be an odd clique. If $d \leq i + l$, then $d \in I$, and the latter never occurs for odd cliques. \square

We remark that the graph $G = G_{d,l}$ defined in the proof of Proposition 2.10 is perfect. Therefore, in the proof we exhibit for each min up/down polytope $P_d(l)$ a corresponding Hansen polytope $\text{Hans}(G_{d,l})$ with equal dimension, number of vertices, and number of facets as $P_d(l)$. It is then natural to wonder if these two polytopes are combinatorially equivalent, or more generally, if min up/down polytopes are just a subclass of Hansen polytopes (after all, both classes are 2-level and centrally symmetric). This turns out not to be the case.

Proposition 2.11. *The min up/down polytope with parameters $d = 8$ and $l = 2$ is not combinatorially equivalent to any Hansen polytope.*

Proof. It can be checked computationally that the min up/down polytope $P_8(2)$ is of dimension 8 and contains 68 vertices, 28 facets, and 604 edges (see Appendix A.1 for details on the computation). The corresponding perfect graph assigned to it in the proof of Proposition 2.10 is P_7 , the path on 7 nodes; and it can be checked as well that its Hansen polytope, $\text{Hans}(P_7)$, is of dimension 8 and contains 68 vertices, 28 facets, and 622 edges (see Appendix A.1). This last number proves that the two polytopes are not combinatorially equivalent.

It remains to show that there is no other perfect graph G , for which $\text{Hans}(G)$ is equivalent to $P_8(2)$. Assume by contradiction that there is such a graph G , with n nodes and m edges, and let \mathcal{C}' and \mathcal{S}' be as in Corollary 2.6. From the information we have on $P_8(2)$, and from the proof of Proposition 2.8, it follows that $n = 7$, $|\mathcal{C}'| = 14$ and $|\mathcal{S}'| = 34$. Notice also that the bound $|\mathcal{C}'| \geq m + n + 1$ gives $m \leq 6$. Suppose first that G is connected; then the bound on m implies that G is a tree. There is extensive bibliography on the number of stable sets on trees, and in particular it is known [82] that $|\mathcal{S}'| \geq F_{n+2}$ (where F_n is the n -th Fibonacci number), and that this bound is tight only in the case of a path. As this bound is tight for G , we conclude that $G = P_7$, a case already considered above.

Now suppose that G is not connected. Then the number $|\mathcal{S}'|$ of stable sets is equal to the product of the corresponding numbers for each connected component. As $|\mathcal{S}'| = 34$ factors into $2 \cdot 17$, G must be composed precisely of two components: an isolated node, and a connected graph G' with $|\mathcal{S}'_{G'}| = 17$ stable sets, $n' = 6$ nodes, and m edges, with $5 \leq m \leq 6$. Now, G' cannot be a tree, as in that case G would only have $|\mathcal{C}'| = 13$ cliques. Therefore, G' must be a *unicyclic graph*, i.e., a tree with an additional edge. There are also extensive results on the number of stable sets on unicyclic graphs; in particular, it is known [97, Thm. 9] that $|\mathcal{S}'_{G'}| \geq F_{n'+1} + F_{n'-1}$. This leads to the inequality $17 \geq 13 + 5$, which is a contradiction. This completes the proof. \square

2.3.4 Polytopes coming from posets

Consider a poset P , with order relation \leq . Its associated *order polytope* is

$$\mathcal{O}(P) = \{x \in [0, 1]^P : x_i \geq x_j \text{ whenever } i \leq j\}, \quad (2.1)$$

and its *chain polytope* is

$$\mathcal{C}(P) = \{x \in \mathbb{R}_+^P : \sum_{i \in I} x_i \leq 1 \text{ for each maximal chain } I \subseteq P\}, \quad (2.2)$$

where we recall that a subset $I \subseteq P$ is a *chain* if every pair of elements in it is comparable. Similarly, $I \subseteq P$ is an *anti-chain* if no pair in it is comparable, and it is a *closed set* if $j \in I$ and $i \leq j$ imply $i \in I$. There is a well-known one-to-one correspondence between the closed sets

and the anti-chains of a poset (the bijection maps a closed set to the subset formed by its maximal elements, which is an anti-chain). Stanley [92] gives the following characterization of vertices of these two polytopes.

Lemma 2.12 ([92]). *The vertices of $\mathcal{O}(P)$ are the characteristic vectors of closed sets of P , and the vertices of $\mathcal{C}(P)$ are the characteristic vectors of the anti-chains of P . In particular, $\mathcal{O}(P)$ and $\mathcal{C}(P)$ have an equal number of vertices.*

From this result it is clear that the order polytope $\mathcal{O}(P)$ is a 2-level polytope because it is a sufficient condition that all vertices have 0/1 coordinates and all facet-defining inequalities can be written as $0 \leq c^T x \leq 1$ for integral vectors c . The chain polytope $\mathcal{C}(P)$ is 2-level as well, as we now explain. Define the comparability graph of P as $G_P([d], E)$, with $\{i, j\} \in E$ whenever $i \leq j$ or $j \leq i$. It is then easy to see that cliques and stable sets of this graph correspond precisely to chains and anti-chains of P , respectively. But as comparability graphs are perfect (see e.g. [17]), it follows that $\mathcal{C}(P)$ is equal to the stable set polytope of G_P , and hence it is 2-level and satisfies Conjecture 2.1 by Proposition 2.7.

The order and chain polytopes of P in general do not have the same number of facets. There is, however, a known relation between these numbers, that immediately gives us our desired bound.

Lemma 2.13 ([55]). *The number of facets of $\mathcal{O}(P)$ is less than or equal to the number of facets of $\mathcal{C}(P)$.*

Lemma 2.14. *Order polytopes and chain polytopes satisfy Conjecture 2.1.*

Proof. Given a poset P on d elements, it is easy to see that both $\mathcal{O}(P)$ and $\mathcal{C}(P)$ are full dimensional, hence both have dimension d . The proof for $\mathcal{C}(P)$ is already given in the lines above. The claimed bound for $\mathcal{O}(P)$ now easily follows from the bounds stated in Lemmas 2.12 and 2.13. If this bound is tight for $\mathcal{O}(P)$, then it must also be tight for $\mathcal{C}(P) = \text{STAB}(G_P)$; this implies by Proposition 2.7 that G_P has no edges, so P is the trivial poset and $\mathcal{O}(P)$ is the cube. \square

To conclude the section, we mention a class of polytopes defined from double posets, which was studied in [14]. A double poset is a triple (P, \leq_+, \leq_-) , where \leq_+ and \leq_- are two partial orders on P . The double order polytope is defined as

$$\mathcal{O}(P, \leq_+, \leq_-) = \text{conv}\{(2\mathcal{O}(P_+) \times \{1\}) \cup (-2\mathcal{O}(P_-) \times \{-1\})\},$$

where P_+ is the poset relative to \leq_+ , and similarly for P_- . A double poset is said to be *compatible* if \leq_+, \leq_- have a common linear extension (i.e. they can be extended to the same total order). In [14] it is proved that, if (P, \leq_+, \leq_-) is compatible, then $\mathcal{O}(P, \leq_+, \leq_-)$ is 2-level if and only if $\leq_+ = \leq_-$ and that in this case the number of its facets is twice the number of chains of (P, \leq_+) . This leads to the following:

Lemma 2.15. *For any poset (P, \leq) , the double order polytope $\mathcal{O}(P, \leq, \leq)$ satisfies Conjecture 2.1.*

Proof. Let $|P| = d$. From the definition, it is clear that $\mathcal{O}(P, \leq, \leq)$ has dimension $d + 1$ and twice as many vertices as $\mathcal{O}(P)$. Let A, C be the sets of anti-chains and chains of P , respectively. Using Lemma 2.12, and the result in [14], we have that $\mathcal{O}(P, \leq, \leq)$ has $2|A|$ vertices and $2|C|$ facets. Now, we remark that Corollary 2.6 applied to the comparability graph of P implies that $|A| \cdot |C| \leq (d + 1)2^d$, this being tight only if P itself is a chain or an anti-chain. The thesis follows immediately. \square

2.3.5 Stable matching polytopes

An instance of the stable matching (or stable marriage) problem, in its most classical version, is defined by a complete bipartite graph $G(M \cup W, E)$ with $n = |M| = |W|$, together with a list $(<_v)_{v \in M \cup W}$, where for each vertex v , $<_v$ is a strict linear order over v 's neighbors. The traditional context of the problem is that there is a set M of men and a set W of women, where each individual wishes to marry a member of the opposite set, and has a list of strict preferences (for instance, $m <_w m'$ means that w prefers m' over m). A *stable marriage* is a perfect matching μ in G with the property that there is no un-matched pair where both individuals prefer each other over their partners; more precisely, if $\mu(v)$ represents v 's partner in matching μ , then μ is stable if and only if

$$\forall mw \in E \setminus \mu, \text{ either } m <_w \mu(w) \text{ or } w <_m \mu(m).$$

Let \mathcal{M} be the set of stable matchings of this instance. The *stable matching polytope* $S(\mathcal{M})$ is the convex hull of the characteristic vectors of all stable matchings in \mathcal{M} . As every instance has at least one stable matching [38], $S(\mathcal{M})$ is a non-empty subset of $[0, 1]^E$. Furthermore, it is known [84] that this polytope can be represented as follows.

$$S(\mathcal{M}) = \left\{ x \in \mathbb{R}_{\geq 0}^E : x(\delta(v)) \leq 1 \ \forall v \in V, \ x_{mw} + \sum_{m' >_w m} x_{m'w} + \sum_{w' >_m w} x_{mw'} \geq 1 \ \forall mw \in E \right\}.$$

From this description, it is evident that $S(\mathcal{M})$ is a 2-level polytope, because all vertices have 0/1 coordinates, and all inequalities are of the form $\alpha \leq c^\top x \leq \alpha + 1$ for some integral vector c and integer α .² Our strategy is to prove that the stable matching polytope is affinely equivalent to an order polytope, and hence satisfies Conjecture 2.1 by Proposition 2.14. To this end, we first present some necessary notation and results.

For a pair of stable matchings μ, μ' in \mathcal{M} , the relation $\mu \preceq \mu'$ signifies that every woman is at least as happy with μ' than with μ , i.e., for each $w \in W$, either $\mu(w) <_w \mu'(w)$ or $\mu(w) = \mu'(w)$. This relation makes \mathcal{M} a distributive lattice; see [66]. We denote by μ_0 and μ_z respectively the

²To visualize this, notice that the above-mentioned description is equivalent to $S(\mathcal{M}) = \{x \in \mathbb{R}^E : 0 \leq x_{mw} \leq 1 \text{ and } 1 \leq x_{mw} + \sum_{m' >_w m} x_{m'w} + \sum_{w' >_m w} x_{mw'} \leq 2 \text{ for each } mw \in E, \text{ and } 0 \leq x(\delta(v)) \leq 1 \text{ for each } v \in V\}$.

(unique) minimum and maximum in this lattice. Further, the ordered pair (μ, μ') of distinct stable matchings is a *covering pair* if $\mu \preceq \mu'$ and there is no other $\mu'' \in \mathcal{M}$ such that $\mu \preceq \mu'' \preceq \mu'$. The lattice structure of \mathcal{M} can be represented by its *Hasse diagram*, which is the directed graph $H(\mathcal{M}, A)$, where A is the set of all covering pairs.

The *rotation* generated by a covering pair $(\mu, \mu') \in A$ is defined as $\rho = (\rho^-, \rho^+)$, where $\rho^- = \mu \setminus \mu'$ and $\rho^+ = \mu' \setminus \mu$. We refer to sets ρ^- and ρ^+ respectively as the tail and the head of rotation ρ .³ Let Π be the set of all rotations generated by covering pairs in A , and notice that more than one covering pair may generate the same rotation in Π . For a pair of rotations ρ, ρ' in Π , we say that ρ *precedes* ρ' , if in any $\mu_0 - \mu_z$ path P in the Hasse diagram H , any arc generating ρ precedes any arc generating ρ' .⁴ This precedence relation defines a poset structure over Π [56]. We now enumerate some properties of the rotation poset Π .

Lemma 2.16. *Let Π be the rotation poset associated to \mathcal{M} .*

1. [52, Thm. 2.5.4] *For each $\mu \in \mathcal{M}$, there is a subset $\Pi(\mu) \subseteq \Pi$ such that, for each $\mu_0 - \mu$ path P in H , the set of rotations generated by arcs in P is precisely $\Pi(\mu)$, with each rotation in it generated exactly once.*
2. [52, Thm. 2.5.7] *For each $\mu \in \mathcal{M}$, $\Pi(\mu)$ is a closed set of the rotation poset Π , and this mapping defines a bijection between \mathcal{M} and the closed sets in Π .*

The following proposition was observed in [29]. We give a proof for completeness.

Proposition 2.17. *The vector family $\{\chi^{\rho^+} - \chi^{\rho^-}\}_{\rho \in \Pi}$ is linearly independent in \mathbb{R}^E .*

Proof. We first prove the following claim: for an edge $mw \in E$ and two rotations $\rho_1, \rho_2 \in \Pi$, if $mw \in \rho_1^+ \cap \rho_2^-$, then ρ_1 precedes ρ_2 . Notice first that ρ_1 and ρ_2 must be in distinct rotations, because the head and the tail of any rotation are always disjoint. Now, consider any $\mu_0 - \mu_z$ path P in H : we know that each of ρ_1 and ρ_2 is generated by an arc in P exactly once, by Lemma 2.16 (1), and we also know that the happiness of woman w increases monotonously along the path. If ρ_2 was generated before ρ_1 , this would imply that w leaves partner m only to go back to him later on, which violates monotonicity. This proves the claim.

To prove the thesis, consider the linear combination

$$\sum_{\rho \in \Pi} \lambda_{\rho} (\chi^{\rho^+} - \chi^{\rho^-}) = 0, \quad (2.3)$$

³This is not the standard definition of rotation found in the literature, but can be seen to be equivalent by [51, Thm. 6]. (Our notation is also different, with the traditional notation being as follows. If $\rho = (\rho^-, \rho^+)$ is generated by (μ, μ') , then ρ is said to be *exposed* in μ ; μ' is said to be obtained from μ after *eliminating* ρ from it, and denoted by μ/ρ ; each edge in ρ^- is *eliminated* by ρ , and each edge in ρ^+ is *produced* by ρ .)

⁴Again, this is not the standard definition of the precedence relation, but can be seen to be equivalent by [52, Thm. 3.2.1].

Chapter 2. On vertices and facets of 2-level polytopes arising in combinatorial settings

for some coefficients λ_ρ , and assume by contradiction that not all coefficients are zero. Among all rotations ρ with $\lambda_\rho \neq 0$, let ρ_2 be a minimal one on the corresponding restriction of the rotation poset, and let mw be an edge in ρ_2^- (such edge exists as no rotation tail can be empty). In [52, Lemma 3.2.1] it is proved that each edge in E appears in the tail of at most one rotation (as well as in the head of at most one rotation in Π). Hence, mw appears in no other tail, so for equation (2.3) to hold, mw must appear in the head of a distinct rotation ρ_1 , with $\lambda_{\rho_1} \neq 0$. By the previous claim, ρ_1 precedes ρ_2 , which contradicts the choice of rotation ρ_2 . This completes the proof. \square

Theorem 2.18. *Given a lattice \mathcal{M} of stable matchings, with associated rotation poset Π , the stable matching polytope $S(\mathcal{M})$ is affinely equivalent to the order polytope $\mathcal{O}(\Pi)$. More precisely, if μ_0 is the minimal element in \mathcal{M} , then*

$$S(\mathcal{M}) = \chi^{\mu_0} + A \cdot \mathcal{O}(\Pi),$$

where $A \in \mathbb{R}^{E \times \Pi}$ is the matrix with columns of the form $A^\rho = \chi^{\rho^+} - \chi^{\rho^-}$ for each $\rho \in \Pi$.

Proof. Let Q be the polytope on the right-hand side of the claimed identity. Q is clearly an affine projection of $\mathcal{O}(\Pi)$ into \mathbb{R}^E . Further, the affine dimension of Q is equal to that of $\mathcal{O}(\Pi)$, by Proposition 2.17. Hence, Q is affinely equivalent to $\mathcal{O}(\Pi)$.

It remains to show that $S(\mathcal{M}) = Q$, which we do by proving that the collection of vertices of these polytopes coincide. Recall from Lemma 2.12 that the vertices of $\mathcal{O}(\Pi)$ are precisely the characteristic vectors of the closed sets in Π , and that these closed sets are in one-to-one correspondence to the stable matchings in \mathcal{M} , by Lemma 2.16 (2). We thus obtain that the vertices of Q are $\{\chi^{\mu_0} + \sum_{\rho \in \Pi(\mu)} (\chi^{\rho^+} - \chi^{\rho^-})\}_{\mu \in \mathcal{M}}$.

Finally, we prove that $\chi^\mu = \chi^{\mu_0} + \sum_{\rho \in \Pi(\mu)} (\chi^{\rho^+} - \chi^{\rho^-})$ for each stable matching μ . Fix $\mu \in \mathcal{M}$, and fix a $\mu_0 - \mu$ path P in H : this defines a chain of stable matchings $\mu_0 \leq \mu_1 \leq \dots \leq \mu_k = \mu$, and a sequence of rotations ρ_1, \dots, ρ_k , so that $\rho_i = (\rho_i^-, \rho_i^+) = (\mu_{i-1} \setminus \mu_i, \mu_i \setminus \mu_{i-1})$ for each $1 \leq i \leq k$. Therefore, $\chi^{\mu_i} = \chi^{\mu_{i-1}} + (\chi^{\rho_i^+} - \chi^{\rho_i^-})$, which by recursion gives us $\chi^\mu = \chi^{\mu_0} + \sum_{i=1}^k (\chi^{\rho_i^+} - \chi^{\rho_i^-})$. By Lemma 2.16 (1), sets $\Pi(\mu)$ and $\{\rho_1, \dots, \rho_k\}$ are equal with no repeated elements. This completes the proof. \square

As remarked before, this result immediately implies our desired bound, by Proposition 2.14.

Corollary 2.19. *Stable matching polytopes satisfy Conjecture 2.1.*

We conclude the section with a remark on Theorem 2.18. Even though its proof is relatively straightforward, to the best of our knowledge this explicit connection was absent in the (extensive) literature of the problem, and it seems to simplify known results as well as shed new light on the structure of the stable matching polytope $S(\mathcal{M})$. In particular, Eirinakis et al. [28] have recently obtained for the first time the dimension, the number of facets,

and a complete minimal linear description of $S(\mathcal{M})$. Their analysis, based on the study of the rotation poset Π , as well as on “reduced non-removable sets of non-stable pairs”, is far from trivial. In contrast, our observation is theoretically simpler and immediately provides those results, as the facial structure of order polytopes is very well understood, and a simple minimal linear description of it is known; see [92]. Moreover, our result is also algorithmically significant, as it provides, from the rotation poset Π , a non-redundant system of equations and inequalities of $S(\mathcal{M})$; and Π can be efficiently constructed from the preference lists, in time $O(n^2)$ [52, Lemma 3.3.2].

2.4 2-Level Matroid Base Polytopes

We start the section with basic definitions and facts about matroids that will be needed throughout the section. For a more complete treatment of these notions we refer the reader to [80]. We identify a matroid M by the couple (E, \mathcal{B}) , where $E = E(M)$ is its ground set, and $\mathcal{B} = \mathcal{B}(M)$ is its base set. Whenever it is convenient, we describe a matroid in terms of its family $\mathcal{I} = \mathcal{I}(M)$ of independent sets or its rank function r_M or simply rk when there is no ambiguity. Given $M = (E, \mathcal{B})$ and a set $F \subseteq E$, the *restriction* $M|F$ is the matroid with ground set F and independent sets $\mathcal{I}(M|F) = \{I \in \mathcal{I}(M) : I \subseteq F\}$; and the *contraction* M/F is the matroid with ground set $M \setminus F$ and rank function $r_{M/F}(A) = r_M(A \cup F) - r_M(F)$. For an element $e \in E$, the *removal of e* is $M - e = M|(E - e)$. A set $F \subseteq E$ is a *circuit* if it is minimally dependent, i.e. F is dependent but every proper subset of it is independent; and $F \subseteq E$ is a *flat* if it is maximal for its rank, i.e. $r(F) < r(F + x)$ for all $x \in E \setminus F$. An element $p \in E$ is called a *loop* (respectively *coloop*) of M if it appears in none (all) of the bases of M .

Consider matroids $M_1 = (E_1, \mathcal{B}_1)$ and $M_2 = (E_2, \mathcal{B}_2)$, with non-empty base sets. If $E_1 \cap E_2 = \emptyset$, we can define the *direct sum* $M_1 \oplus M_2$ as the matroid with ground set $E_1 \cup E_2$ and base set $\mathcal{B}_1 \times \mathcal{B}_2$. If, instead, $E_1 \cap E_2 = \{p\}$, where p is neither a loop nor a coloop in M_1 or M_2 , we let the *2-sum* $M_1 \oplus_2 M_2$ be the matroid with ground set $E_1 \cup E_2 - p$, and base set $\{B_1 \cup B_2 - p : B_i \in \mathcal{B}_i \text{ for } i = 1, 2 \text{ and } p \in B_1 \Delta B_2\}$. A matroid is *connected* (2-connected for some authors) if it cannot be written as the direct sum of two matroids, each with fewer elements; and a connected matroid M is *3-connected* if it cannot be written as a 2-sum of two matroids, both with strictly fewer elements than M .

The proofs of the following facts can be found e.g. in [80].

Proposition 2.20. *Let $M = M_1 \oplus_2 M_2$, with $E(M_1) \cap E(M_2) = \{p\}$.*

1. $M_1 \oplus_2 M_2$ is connected if and only if so are M_1 and M_2 .
2. $\mathcal{B}(M_1 \oplus_2 M_2) = \mathcal{B}(M_1 - p) \times \mathcal{B}(M_2 / p) \uplus \mathcal{B}(M_1 / p) \times \mathcal{B}(M_2 - p)$.
3. $|\mathcal{B}(M_i)| = |\mathcal{B}(M_i - p)| + |\mathcal{B}(M_i / p)|$, for $i = 1, 2$.
4. If $M_2 = M'_2 \oplus M''_2$, where $E(M_1) \cap E(M'_2) = \{p\}$ and $(E(M_1) \cup E(M'_2)) \cap E(M''_2) = \emptyset$, then $M_1 \oplus_2 M_2 = (M_1 \oplus_2 M'_2) \oplus M''_2$.

2.4.1 2-level matroid polytopes and Conjecture 2.1

In this section we describe those matroids whose base polytope is 2-level and we prove that Conjecture 2.1 holds for such polytopes.

The *base polytope* $B(M) \subseteq \mathbb{R}^E$ of a matroid $M = (E, \mathcal{B})$ (also called matroid polytope) is given by the convex hull of the characteristic vectors of its bases. The following is known to be a description of $B(M)$ (see, for instance, [89]):

$$B(M) = \{x \in [0, 1]^E : x(F) \leq r(F) \text{ for } F \subseteq E; x(E) = r(E)\}. \quad (2.4)$$

A matroid $M(E, \mathcal{B})$ is *uniform* if $\mathcal{B} = \binom{E}{k}$, where k is the rank of M . We denote the uniform matroid with n elements and rank k by $U_{n,k}$. It is easy to check that the base polytope of a uniform matroid is a hypersimplex, i.e. $B(U_{n,k}) = \{x \in \mathbb{R}^n : 0 \leq x_i \leq 1, \sum_{i=1}^n x_i = k\}$. Notice that, if M_1 and M_2 are uniform matroids with $|E(M_1) \cap E(M_2)| = 1$, then $M_1 \oplus_2 M_2$ is unique up to isomorphism, for any possible common element.




Figure 2.1 – A representation of $M = U_{5,2} \oplus_2 U_{6,3}$. M has ground set $\{1, 2, 3, 4, 6, 7, 8, 9, 10\}$ and rank 4, and two of its bases are $\{1, 2, 6, 7\}$ and $\{1, 6, 7, 8\}$. $B(M)$ is 2-level (see Theorem 2.21).

Let \mathcal{M} be the class of matroids whose base polytope is 2-level. \mathcal{M} has been characterized in [48]:

Theorem 2.21. *The base polytope of a matroid M is 2-level (i.e. $M \in \mathcal{M}$) if and only if M can be obtained from uniform matroids through a sequence of direct sums and 2-sums.*

The following lemma implies that we can, when looking at matroids in \mathcal{M} , decouple the operations of 2-sum and direct sum.

Lemma 2.22. *Let M be a matroid obtained by applying a sequence of direct sums and 2-sums from the matroids M_1, \dots, M_k . Then $M = M'_1 \oplus M'_2 \oplus \dots \oplus M'_t$, where each of the M'_i is obtained by repeated 2-sums from some of the matroids M_1, \dots, M_k .*

Proof. Immediately from repeated applications of Proposition 2.20, part 4. □

It is immediate to see that if $M = M_1 \oplus M_2$, then $B(M)$ is equal to the Cartesian product $B(M_1) \times B(M_2)$. This, together with Lemma 2.22, suggests that when investigating matroids in \mathcal{M} , the interesting case is when such matroids are connected.

Proposition 2.23. *Let $M \in \mathcal{M}$ be connected and non-uniform, with $M = U_1 \oplus_2 \dots \oplus_2 U_t$, where U_i are uniform matroids and $t > 1$. Then we can assume without loss of generality that every U_i has at least 3 elements.*

Proof. No matroid in a 2-sum can have ground set of size one, since the 2-sum is defined when the common element is not a loop or a coloop of either summand. For the same reason, we can exclude the matroids $U_{2,0}, U_{2,2}$. The only remaining uniform matroid on two elements is $U_{2,1}$. However, it is easy to see that for any matroid M , $M \oplus_2 U_{2,1}$ is isomorphic to M : if the ground set of $U_{2,1}$ is $\{p, e\}$, with p being the element common to M , the 2-sum has the only effect of replacing p by e in M . \square

We now make a general observation on the structure of the base polytope of 2-sums of matroids, which will be used to prove all the results in this section. This fact can be derived from [48], Lemma 3.4, and a weaker version of it is also observed in [58], but for completeness we give a simple proof.

Lemma 2.24. *Let $M_1(E_1, \mathcal{B}_1), M_2(E_2, \mathcal{B}_2)$ be matroids with $E_1 \cap E_2 = \{p\}$ and let $M = M_1 \oplus_2 M_2$. Then $B(M)$ is linearly isomorphic to $B(M_1) \times B(M_2) \cap \{x \in \mathbb{R}^{E_1 \uplus E_2} : x_{p_1} + x_{p_2} = 1\}$, where $E_1 \uplus E_2 = E_1 \cup E_2 \cup \{p_1, p_2\} - p$ is the disjoint union of E_1 and E_2 , with p_1 and p_2 corresponding to $p \in E_1$ and $p \in E_2$ respectively.*

Proof. Let $Q = B(M_1) \times B(M_2) \cap H$, where $H = \{x \in \mathbb{R}^{E_1 \uplus E_2} : x_{p_1} + x_{p_2} = 1\}$, let $E = E_1 \cup E_2 - p$ be the ground set of M , and consider the projection $\varphi : \mathbb{R}^{E_1 \uplus E_2} \rightarrow \mathbb{R}^E$, i.e. such that $\varphi(\vec{1}_e) = \vec{1}_e$ for any $e \in E$, and $\varphi(\vec{1}_{p_1}) = \varphi(\vec{1}_{p_2}) = \vec{0}$. We first claim that $B(M) = \varphi(Q)$. It follows directly from definitions of cartesian product and 2-sum that $B(M)$ and $\varphi(Q)$ have the same integer vertices. To prove the claim we need to show that Q (hence $\varphi(Q)$) does not have any fractional vertices. Suppose that such a vertex v exists: then v is the intersection of the hyperplane H with (the interior of) an edge of $B(M_1) \times B(M_2)$. Using the properties of adjacency of the cartesian product, we can assume without loss of generality that $v = \lambda w + (1 - \lambda) w'$ for some $0 < \lambda < 1$, where $w = (\chi^{B_1}, \chi^{B_2}), w' = (\chi^{B'_1}, \chi^{B'_2})$ are vertices of $B(M_1) \times B(M_2)$, with $\chi^{B_2}, \chi^{B'_2}$ adjacent vertices of $B(M_2)$. In particular we have that $w_{p_1} = w'_{p_1}$, hence $w_{p_1} + w_{p_2} \geq 1, w'_{p_1} + w'_{p_2} \geq 1$, but this is a contradiction since w, w' must be on two different sides of H .

We are left to show that φ restricted to Q is injective to conclude that φ is a bijection from Q to $B(M)$. To see this, assume that there are $x, y \in Q$ such that $\varphi(x) = \varphi(y)$, hence $x_e = y_e$ for any $e \in E$. But then since x, y satisfy the rank equality of $B(M_1)$,

$$x_{p_1} = \text{rk}(M_1) - \sum_{e \in E_1 - p} x_e = \text{rk}(M_1) - \sum_{e \in E_1 - p} y_e = y_{p_1},$$

and arguing similarly we get $x_{p_2} = y_{p_2}$, therefore we have $x = y$. \square

Before we prove that Conjecture 2.1 holds for 2-level base polytopes, we need one last technical

ingredient which is a consequence of the previous Lemma.

Proposition 2.25. *Let $M \in \mathcal{M}$ be such that $M = M_1 \oplus_2 U$ where $U = U_{n,k}$ is a 3-connected uniform matroid with $n \geq 3$. Then $f_{d-1}(B(M)) \leq f_{d_1-1}(B(M_1)) + 2(n-1)$, where d denotes the dimension of $B(M)$; and if $n = 3$ then $f_{d-1}(B(M)) \leq f_{d_1-1}(B(M_1)) + 2$.*

Proof. Using Lemma 2.24, we obtain that $B(M)$ is linearly isomorphic to $Q = B(M_1) \times B(U) \cap \{x \in \mathbb{R}^{E_1 \sqcup E_2} : x_{p_1} + x_{p_2} = 1\}$, where E_1, E_2, p, p_1, p_2 are defined as before. From this it follows that $f_{d-1}(B(M)) \leq f_{d_1-1}(B(M_1)) + f_{d_2-1}(B(U))$, where d_2 is the dimension of $B(U)$. Moreover, as already remarked $B(U) = \{x \in \mathbb{R}^{d_2} : 0 \leq x \leq 1, \sum_i x_i = k\}$ hence $f_{d_2-1}(B(U)) \leq 2n$ and $f_{d-1}(B(M)) \leq f_{d_1-1}(B(M_1)) + 2n$. To slightly sharpen the bound, we claim that the inequalities $0 \leq x_{p_2} \leq 1$ present in the description of Q are redundant, which proves the first part of the thesis. Indeed, they are immediately implied by the inequalities $0 \leq x_{p_1} \leq 1$ (which must be implied by the description of $B(M_1)$) together with the equation $x_{p_1} + x_{p_2} = 1$.

We now consider the case $n = 3$. It is immediate to check that there are two cases, $U = U_{3,1}$ and $U = U_{3,2}$, but for both $B(U)$ is isomorphic to a triangle in the plane, and hence $f_{d_2-1}(B(U)) = 3$, with one inequality for each variable: for instance, a description of $B(U_{3,1})$ is $\{x \in \mathbb{R}^3 : x \geq 0, x_1 + x_2 + x_3 = 1\}$. Arguing as before, we obtain that in the resulting description of Q the inequality relative to x_{p_2} is redundant, thus getting the desired bound. \square

Theorem 2.26. *2-level matroid base polytopes satisfy Conjecture 2.1.*

Proof. We will use the fact that, for any $n \geq 3$ and any $k \in \{0, \dots, n\}$, $\binom{n}{k} \leq \frac{3}{4}2^{n-1}$. This can be easily proved by induction. We prove the conjecture on the polytope $B(M)$, for each matroid $M = (E, \mathcal{B}) \in \mathcal{M}$, and we prove it by induction on the number of elements $n = |E|$. The base cases $n \leq 3$ can be easily verified.

If M is not connected, then $M = M_1 \oplus M_2$ for two matroids $M_1, M_2 \in \mathcal{M}$, each with fewer elements than M , so by induction hypothesis the conjecture holds for them. As already remarked, the base polytope $B(M)$ is simply the Cartesian product $B(M_1) \times B(M_2)$, so by Lemma 2.3 the conjecture also holds for $B(M)$, and is tight only if $B(M)$ is a cube.

Assume from now on that M is connected. In [48], it is proven that the smallest affine subspace containing the base polytope of a connected matroid on n elements is of dimension $d = n - 1$. If M is uniform, $M = U_{n,k}$, the number of vertices in $B(M)$ is $f_0 = |\mathcal{B}| = \binom{n}{k} \leq \frac{3}{4}2^{n-1}$, where we assumed $n \geq 3$. And in view of Proposition 2.23, the constraints of the form $0 \leq x \leq 1$ are sufficient to define $B(M)$, hence the number of facets is $f_{d-1} \leq 2n$. Therefore, $f_0 f_{d-1} \leq \frac{3}{4}n2^n \leq (n-1)2^n = d2^{d+1}$, where the last inequality is loose for $n \geq 5$. The only examples with $n \leq 4$ for which the conjecture is tight correspond to cubes, and the 3-dimensional cross-polytope coming from $U_{4,2}$.

Finally, assume that M is connected but is not uniform, so it is not 3-connected. Then $M = M_1 \oplus_2 M_2$, with matroids $M_1, M_2 \in \mathcal{M}$ each with fewer elements than M , so by induction

hypothesis the conjecture holds for both of them. Let $E(M_1) \cap E(M_2) = \{p\}$. Both M_1 and M_2 are connected, by Proposition 2.20. We can assume without loss of generality that $E(M_1) = n_1 \geq n_2 = E(M_2)$, and that M_2 is uniform, $M_2 = U_{n_2, k_2}$, with $n_2 \geq 3$ (by Proposition 2.23). We consider two cases for the value of n_2 .

Case $n_2 \geq 4$: first notice that the family \mathcal{M} is closed under removing or contracting an element. This is because if $e \in M \in \mathcal{M}$, the base polytopes $B(M - e)$ and $B(M / e)$ are affinely isomorphic to the faces of $B(M)$ that intersect the hyperplanes $x_e = 0$ and $x_e = 1$, respectively, and by Lemma 2.2 these faces are also 2-level. Hence, we know from Proposition 2.20 that

$$\begin{aligned} f_0 &= |\mathcal{B}(M)| = |\mathcal{B}_{M_1-p}| \cdot |\mathcal{B}_{U_{n_2, k_2}/p}| + |\mathcal{B}_{M_1/p}| \cdot |\mathcal{B}_{U_{n_2, k_2}-p}| \\ &= \binom{n_2-1}{k_2-1} |\mathcal{B}_{M_1-p}| + \binom{n_2-1}{k_2} |\mathcal{B}_{M_1/p}| \\ &\leq \frac{3}{4} 2^{n_2-2} (|\mathcal{B}_{M_1-p}| + |\mathcal{B}_{M_1/p}|) = \frac{3}{4} 2^{d_2-1} |\mathcal{B}(M_1)|. \end{aligned}$$

From Proposition 2.25, the number of facets in $B(M)$ is $f_{d-1}(B(M)) \leq f_{d-1}(B(M_1)) + 2(n_2 - 1) = f_{d-1}(B(M_1)) + 2d_2$. We use the induction hypothesis in M_1 , and the trivial bound $|\mathcal{B}(M_1)| \leq 2^{d_1}$ to obtain:

$$\begin{aligned} f_0 f_{d-1}(B(M)) &< \frac{3}{4} 2^{d_2-1} |\mathcal{B}(M_1)| (f_{d-1}(B(M_1)) + 2d_2) \\ &\leq \frac{3}{4} 2^{d_2-1} (d_1 2^{d_1+1} + 2^{d_1} (2d_2)) \\ &= \frac{3}{4} (d_1 + d_2) 2^{d_1+d_2} < (d_1 + d_2 - 1) 2^{d_1+d_2} = d 2^{d+1}. \end{aligned}$$

Where in the last inequality we used the fact that $n_1 \geq n_2 \geq 4$, so $d_1 \geq d_2 \geq 3$.

Case $n_2 = 3$: We can prove in a similar manner as before that

$$f_0 = |\mathcal{B}(M)| < \binom{2}{1} (|\mathcal{B}(M_1 - p)| + |\mathcal{B}(M_1 / p)|) = 2|\mathcal{B}(M_1)|.$$

And from Proposition 2.25, $f_{d-1}(B(M)) \leq f_{d-1}(B(M_1)) + 2$. Thus,

$$f_0 f_{d-1}(B(M)) < 2|\mathcal{B}(M_1)| (f_{d-1}(B(M_1)) + 2) \leq 2(d_1 2^{d_1+1} + 2^{d_1} \cdot 2) = d 2^{d+1}.$$

We conclude by remarking that, since the inequalities above hold strictly, the only 2-level base polytopes satisfying the bound of Conjecture 2.1 are cubes and cross-polytopes. \square

As the forest matroid of a graph G is in \mathcal{M} if and only if G is series-parallel [48], we deduce the following.

Corollary 2.27. *Conjecture 2.1 is true for the spanning tree polytope of series-parallel graphs.*

Throughout the rest of the section, we assume that $M_1(E_1, \mathcal{B}_1)$, $M_2(E_2, \mathcal{B}_2)$ are connected matroids, with, $E_1 \cap E_2 = \{p\}$, and we define $M = M_1 \oplus_2 M_2$. Thanks to Proposition 2.20, under these assumptions M is also connected. By the arguments above, characterizing $B(M)$ essentially boils down to characterizing flacets of $M_1 \oplus_2 M_2$.

To prove Theorem 2.30 we first need the following technical observation.

Observation 2.29. *Let $M_1(E_1, \mathcal{B}_1)$, $M_2(E_2, \mathcal{B}_2)$, $M(E, \mathcal{B}) = M_1 \oplus_2 M_2$ be as above. Let $\emptyset \neq F \subseteq E_i - p$ for some $i \in \{1, 2\}$, and assume that p is not a loop or a coloop of $M_i|(F + p)$. Then, for $j \in \{1, 2\}$, $i \neq j$, one has*

$$M|(E_j \cup F - p) = M_i|(F + p) \oplus_2 M_j,$$

hence $\text{rk}(E_j \cup F - p) = \text{rk}(E_j - p) + \text{rk}(F) - 1$.

Proof. The two matroids clearly have the same ground set. We will show that they also have the same independent sets. For simplicity fix $i = 1$, $j = 2$, and let $I \subseteq E_2 \cup F - p$. Then, as M is a 2-sum, I is an independent set of $M|(E_2 \cup F - p)$ if and only if $I = I_1 \cup I_2 - p$, for some I_1, I_2 independent sets of M_1, M_2 respectively with $p \in I_1 \triangle I_2$, where in particular I_1 can be chosen to be a subset of $F + p$. Hence the latter is equivalent to I being an independent set of $M_1|(F + p) \oplus_2 M_2$. \square

Theorem 2.30. *Let $M_1(E_1, \mathcal{B}_1)$, $M_2(E_2, \mathcal{B}_2)$, $M(E, \mathcal{B}) = M_1 \oplus_2 M_2$ be as above. Let $F \subsetneq E$. Then F is a flacet of M if and only if one of the following holds:*

1. $F = E_i \cup F' - p$, where F' is a flacet of M_j containing p , and $i \neq j \in \{1, 2\}$.
2. F is a flacet of M_i not containing p for some $i \in \{1, 2\}$.
3. $F = E_i - p$, and $M_i - p$, M_j / p are connected, for some $i, j \in \{1, 2\}$ with $i \neq j$.

Proof. From Lemma 2.24, using the same notation, we have that $B(M) = \Pi(Q)$, where $Q = B(M_1) \times B(M_2) \cap \{y \in \mathbb{R}^{E_1 \sqcup E_2} : y_{p_1} + y_{p_2} = 1\}$ and $\Pi : \mathbb{R}^{E_1 \sqcup E_2} \rightarrow \mathbb{R}^E$ is the function that projects out the components corresponding to p_1, p_2 , and it is a bijection from Q to $B(M)$. Let rk_i denote the rank function of M_i , for $i = 1, 2$. Let us consider the following description of Q , obtained by adding together a non-redundant description of $B(M_1)$ and $B(M_2)$:

$$\begin{aligned} Q = \{y \in \mathbb{R}^{E_1 \sqcup E_2} : & \quad y_e \geq 0 & \quad \text{for } e \in E_i : M_i - e \text{ is connected, } i = 1, 2 \\ & y(F) \leq \text{rk}_i(F) & \quad \text{for } F \text{ flacet of } M_i, i = 1, 2 \\ & y(E_i) = \text{rk}_i(E_i) & \quad \text{for } i = 1, 2 \\ & y_{p_1} + y_{p_2} = 1\}. \end{aligned}$$

We argue that this is a non-redundant description of Q , possibly apart from the inequalities $y_{p_i} \leq 1$. First, notice that the description without the equation $y_{p_1} + y_{p_2} = 1$ is a non-redundant

description of $B(M_1) \times B(M_2)$. By adding such equation, the inequalities $y_{p_i} \leq 1$ and $y_{p_j} \geq 0$ become equivalent for $i \neq j$, hence, in case both of them were present, one becomes redundant; however it is easy to see no other inequality can become redundant this way.

Now, fix $\emptyset \neq F \subsetneq E$ and let $\mathcal{F} = \{x \in B(M) : x(F) = \text{rk}(F)\}$ be a face of $B(M)$. We have that \mathcal{F} is a facet of $B(M)$ if and only if $\mathcal{F}' = \Pi^{-1}(\mathcal{F})$ is a facet of Q . It is immediate to see that the non-negativity facets $\{x \in B(M) : x_e = 0\}$ are in one to one correspondence with the facets $\{y \in Q : y_e = 0\}$ for $e \in E = E_1 \uplus E_2 - \{p_1, p_2\}$. This implies that facets of M must be in correspondence with the facets of M_1, M_2 , and possibly the inequalities $y_{p_1} \geq 0, y_{p_2} \geq 0$ (or equivalently $y_{p_1} \leq 1, y_{p_2} \leq 1$). Hence, we can characterize F by investigating the inequality corresponding to \mathcal{F}' , using the fact that each facet of a base polytope has a unique representation in term of non-negativity or facet inequality. We have that F is a facet if and only if one of the following holds:

1. $\mathcal{F}' = \{y \in Q : y(F') = \text{rk}_i(F')\}$ for some F' facet of M_i , $i \in \{1, 2\}$, with $p \notin F'$. Then clearly $\mathcal{F} = \Pi(\mathcal{F}') = \{x \in B(M) : x(F') = \text{rk}(F')\}$, equivalently $F = F'$.
2. $\mathcal{F}' = \{y \in Q : y(F') = \text{rk}_i(F')\}$ for some F' facet of M_i , $i \in \{1, 2\}$, with $p \in F'$. Fix $i = 1$ for simplicity. We now argue that, for any basis of M $B = B_1 \cup B_2 - p$, we have that $|B_1 \cap F'| = \text{rk}_1(F')$ if and only if $|B \cap (E_2 \cup F' - p)| = \text{rk}(E_2 \cup F' - p)$. Indeed, using Observation 2.29, we have

$$|B \cap (E_2 \cup F' - p)| = |B_1 \cap F'| + \text{rk}_2(E_2) - 1 = \text{rk}(E_2 \cup F' - p) \iff |B_1 \cap F'| = \text{rk}_1(F').$$

Hence, this case is equivalent to $F = E_2 \cup F' - p$.

3. $\mathcal{F}' = \{y \in Q : y_{p_i} = 0\}$ for some $i \in \{1, 2\}$. Fix $i = 1$. Notice that for any basis of M $B = B_1 \cup B_2 - p$, we have $|B \cap (E_1 - p)| = \text{rk}(E_1 - p)$ if and only if $p \notin B_1$, hence $\mathcal{F} = \Pi(\mathcal{F}') = \{x \in B(M) : x(E_1 - p) = \text{rk}(E_1 - p)\}$, equivalently $F = E_1 - p$. We still need to argue that this case happens if and only if $M_1 - p$ and M_2/p are connected. The ‘if’ direction is clear. For the other direction, if \mathcal{F}' is a facet of Q , then one of $y_{p_1} \geq 0, y_{p_2} \leq 1$ is included in the non-redundant description of Q given above. Assume it is $y_{p_1} \geq 0$, which implies that $M_1 - p$ is connected. If M_2/p was not connected, $y_{p_2} \leq 1$ would be a redundant inequality in the description of $B(M_2)$ and it would be implied by some of the others. But then $y_{p_1} \geq 0$ would be implied as well by the same inequalities plus $y_{p_1} + y_{p_2} = 1$, a contradiction. One can argue similarly in the case $y_{p_2} \leq 1$ was included in the description of Q .

□

Corollary 2.31. *The following is a non-redundant description of $B(M)$:*

$$\begin{aligned}
 B(M) = \{x \in \mathbb{R}^E : \\
 \begin{array}{lll}
 x_e & \geq 0 & e \in E_i - p : M_i - e \text{ connected}, i = 1, 2 \\
 x(E_i \cup F - p) & \leq r(E_i \cup F - p) & F \text{ flacet of } M_j : \{p\} \subsetneq F, i \neq j \in \{1, 2\} \\
 x(F) & \leq r(F) & F \text{ flacet of } M_i : p \notin F, i \in \{1, 2\} \\
 x(E_i - p) & \leq r(E_i - p) & \text{if } M_i - p, M_j / p \text{ connected}, i \neq j \in \{1, 2\} \\
 x(E) & = r(E).
 \end{array}
 \end{aligned} \tag{2.7}$$

2.4.3 Linear Description of 2-Level Matroid Base Polytopes

Using Proposition 2.25, one can easily prove by induction that for any $M \in \mathcal{M}$ the number of facets of $B(M)$ is linear in the size of the ground set. However, the description of $B(M)$ given in (2.5) has exponentially many inequalities. Finding compact description for the base and the independent set polytopes of matroids has been the object of many studies, especially in terms of extended formulations: see [85] for a negative result, and [19], [57] for formulations for special classes of matroids. These results can be seen as generalizations of the formulations given for the spanning tree polytope by Martin [78]. In this section we give an explicit description of 2-level base matroids with linearly many inequalities. The rank inequalities needed in our description have a natural interpretation in terms of the combinatorial structure of the matroid, and our description can be obtained in polynomial time in the size of the ground set of the matroid (given an independence oracle for it).

In light of Theorem 2.21, to obtain a linear description of $B(M)$ for $M \in \mathcal{M}$ one needs to investigate the base polytope of uniform matroids, and how the description of base polytopes behaves with respect to the operations of 1-sum and 2-sum of matroids. For the former, one can make the following easy observation, which has already been stated in equivalent form in Section 2.4.1:

Observation 2.32. *Let $U_{n,k}$ be a uniform matroid. Then $U_{n,k}$ has no facet beside its singletons. In particular, $B(U_{n,k}) = \{x \in \mathbb{R}^n : 0 \leq x \leq 1, \sum_i x_i = k\}$.*

Since the base polytope of the 1-sum of matroids is the Cartesian product of the base polytopes, to obtain a linear description of $B(M)$ for $M \in \mathcal{M}$, we can focus on base polytopes of connected matroids. We will use Theorem 2.30 to deal with the base polytope of 2-sums of matroids. Any connected matroid can be seen as a sequence of 2-sums, which can be represented via a tree (see Figure 2.2): the following is a version of [80, Proposition 8.3.5] tailored to our needs. For completeness, we include a proof.

Theorem 2.33. *Let M be a connected matroid. Then there are 3-connected matroids M_1, \dots, M_t , and a t -vertex tree $T = T(M)$ with edges labeled e_1, \dots, e_{t-1} and vertices labeled M_1, \dots, M_t , such that*

1. $E(M) \cap \{e_1, \dots, e_{t-1}\} = \emptyset$, and $E(M_1) \cup E(M_2) \cup \dots \cup E(M_t) = E(M) \cup \{e_1, \dots, e_{t-1}\}$;

2. if the edge e_i joins the vertices M_{j_1} and M_{j_2} , then $E(M_{j_1}) \cap E(M_{j_2}) = \{e_i\}$;
3. if no edge joins the vertices M_{j_1} and M_{j_2} , then $E(M_{j_1}) \cap E(M_{j_2}) = \emptyset$.

Moreover, M is the matroid that labels the single vertex of the tree $T/e_1, \dots, e_{t-1}$ at the conclusion of the following process: contract the edges e_1, \dots, e_{t-1} of T one by one in order; when e_i is contracted, its ends are identified and the vertex formed by this identification is labeled by the 2-sum of the matroids that previously labeled the ends of e_i .

Proof. We proceed by induction on $n = |E(M)|$. For $n = 1$, M is 3-connected, T consists of only one vertex and there is nothing to show. For $n > 1$: if M is 3-connected, again there is nothing to show. Otherwise, $M = M' \oplus_2 M''$ for some matroids M', M'' , that are connected (due to Proposition 2.20) and that satisfy $|E(M')|, |E(M'')| < n$. Hence by induction hypothesis the thesis holds for M', M'' . Let T', T'' be their corresponding trees, with vertices labeled by the 3-connected matroids M'_1, \dots, M'_{t_1} and M''_1, \dots, M''_{t_2} respectively, edges labeled e'_1, \dots, e'_{t_1-1} and $e''_1, \dots, e''_{t_2-1}$ respectively, and let $t = t_1 + t_2$. By definition of 2-sum there is exactly one element, which we denote by e_{t-1} , in $E(M') \cap E(M'')$. By induction we have:

$$\begin{aligned} E(M) &= E(M') \cup E(M'') \setminus \{e_{t-1}\} \\ &= \\ &= (E(M'_1) \cup \dots \cup E(M'_{t_1}) \setminus \{e'_1, \dots, e'_{t_1-1}\}) \cup (E(M''_1) \cup \dots \cup E(M''_{t_2}) \setminus \{e''_1, \dots, e''_{t_2-1}\}) \setminus \{e_{t-1}\}. \end{aligned}$$

We can assume without loss of generality that $\{e'_1, \dots, e'_{t_1-1}\} \cap E(M'') = \emptyset$ by renaming the elements of $E(M'')$, and similarly we can assume $\{e''_1, \dots, e''_{t_2-1}\} \cap E(M') = \emptyset$. Since M' satisfies properties 1-3, there is exactly one matroid M'_i such that $e_{t-1} \in E(M'_i)$, and similarly there is exactly one matroid M''_j such that $e_{t-1} \in E(M''_j)$. Let T be the tree obtained by joining T', T'' through the edge (M'_i, M''_j) . Now, it is easy to check that the matroids labeling the vertices of T will satisfy properties 1-3 after an appropriate renaming of the matroids and relabeling of the edges (M'_i will be renamed M_i , M''_j M_{j+t_1} , and similarly for the elements e'_i, e''_j). The statement about the contraction $T/e_1, \dots, e_{t-1}$ follows by induction: one first contracts the edges in T' (e_1, \dots, e_{t_1-1}), then the edges in T'' (e_{t_1}, \dots, e_{t-2}), obtaining vertices labeled by M' and M'' . Then, contracting the edge e_{t-1} joining M', M'' one gets $M' \oplus_2 M'' = M$. \square

Example 2.34. Consider the matroid M whose associated tree structure is given in Figure 2.2. The ground set of M is $\{1, 2, 3, 4, 8, 9, 10, 11, 12, 13, 14, 15\}$ and its rank, which can be computed as the sum of the ranks of the nodes minus the number of edges, is 4. $\{1, 2, 11, 13\}$ is a basis.

For a connected matroid $M(E, \mathcal{B}) \in \mathcal{M}$, Theorem 2.33 reveals a tree structure $T(M)$, where every node represents a 3-connected uniform matroid, and every edge represents a 2-sum operation. We now give a simple description of the associated base polytope. Let a be an edge of $T(M)$. The removal of a breaks T into two connected components C_a^1 and C_a^2 . Let E_a^1




Figure 2.2 – The matroid from Example 2.34.

(resp. E_a^2) be the set of elements from E that belong to uniform matroids from C_a^1 (resp. C_a^2). The following theorem shows that the inequalities needed to describe $B(M)$ are the “trivial” inequalities $0 \leq x \leq 1$, plus $x(F) \leq r(F)$, where $F = E_a^1$ or E_a^2 for some edge a of $T(M)$. If M is 2-sum of uniform matroids U_1, \dots, U_t , then clearly T will have $t - 1$ edges. From Proposition 2.23, we know that $E(U_i) \geq 3$ for any i . Hence, if $|E| = n$, we have

$$n = \sum_{i=1}^t |E(U_i)| - 2(t-1) \geq 3t - 2(t-1) = t + 2,$$

hence $t \leq n - 2$. Thus, the total number of inequalities needed is linear in the number of elements.

Theorem 2.35. *Let $M = (E, \mathcal{B}) \in \mathcal{M}$ be a connected matroid obtained as 2-sum of uniform matroids $U_1 = U_{n_1, k_1}, \dots, U_t = U_{n_t, k_t}$. Let $T(N, A)$ be the tree structure of M according to Theorem 2.33. For each $a \in A$, let $C_a^1, C_a^2, E_a^1, E_a^2$ be defined as above. Then*

$$\begin{aligned} B(M) = \{x \in \mathbb{R}^E : & \quad x \geq 0 \\ & \quad x \leq 1 \\ & \quad x(F) \leq \text{rk}(F) \quad \text{for } F = E_a^i \text{ for some } i \in \{1, 2\} \text{ and } a \in A, \\ & \quad x(E) = \text{rk}(E)\}. \end{aligned}$$

Moreover, if $F = E_a^i$ for some $i \in \{1, 2\}$ and $a \in A$, then $r(F) = 1 - |C_a^i| + \sum_{j: U_j \in C_a^i} k_j$.

Proof. Let us call a subset $C \subseteq N$ a *valid component* for T if $C = C_a^i$ for some $i \in \{1, 2\}$ and $a \in A$, and denote the set of all valid components of T by \mathcal{F} . Each connected subtree of $T(N, A)$ represents a connected matroid obtained as 2-sums of uniform matroids. Thus, we can prove the theorem by induction on t . The statement on the rank is immediate. For $t = 1$, \mathcal{F} is empty and, thanks to Observation 2.32, the remaining inequalities are enough to describe $B(M)$. Now let $t > 1$. Thanks to Theorem 2.30, to prove the thesis it is enough to show that, if F is a flacet of M with $|F| \geq 2$, then $F \in \mathcal{F}$. First notice that we can write, without loss of

generality, $M = M' \oplus_2 U_t$, where U_t corresponds to a leaf v_t of T and M' is obtained as 2-sums of U_1, \dots, U_{t-1} , hence it satisfies the inductive hypothesis. Note that the tree corresponding to M' is then $T - v_t$. Let us denote by v_l the only neighbor of v_t in T . Let $E' + p$, $E(U_t) = E_t + p$ be the ground sets of M' , U_t respectively, where $E' = \bigcup_{i=1}^{t-1} E_i$, and $E_i = E \cap E(U_i)$ for $i = 1, \dots, t$. Clearly $p \in E(U_l)$. Now, since F is a facet of M , we can apply Theorem 2.30 to get three possible cases. If F has non-empty intersection with both $E(M')$ and E_t , then we are in case 1 and either $F = E(U_t) \cup F' - p$ or $F = E' \cup F_t - p$, where F', F_t are facets of M' , U_t respectively, containing p . However, the latter case is not possible because of Observation 2.32, so the only possibility is that $F = E_t \cup F'$. By induction, F' belongs to \mathcal{F}' defined for M' as in the statement of the theorem. Moreover, since F' contains p , its corresponding component C in $T - v_t$ contains v_l and then $C + v_t$ is a valid component for T . Moreover $|F' \cap E_i| \in \{0, |E_i|\}$ for any $i = 1, \dots, t-1$, which implies $F \in \mathcal{F}$. Suppose now we are in case 2, i.e., F is strictly contained in one of E', E_t . Then F is a facet of one of M' , U_t , the latter not being possible again due to Observation 2.32. So F is a facet of M' and it does not contain p , hence by induction hypothesis its corresponding component C does not contain v_l . But then C is a valid component of T and again $F \in \mathcal{F}$. Finally, if we are in case 3 then $F = E_t$ or $F = E$, and in both cases $F \in \mathcal{F}$. \square

We conclude by remarking that, for any matroid M , the corresponding tree structure given in Theorem 2.33 can be obtained in polynomial time, given an independence oracle for M , for instance using the shifting algorithm given in [9]. This means that, given an independence oracle for $M \in \mathcal{M}$, one can efficiently write down the description of $B(M)$ given by Theorem 2.35: first, one obtains the tree structure and the corresponding uniform matroids, and then the rank inequalities corresponding to the edges of the tree. The latter part just takes linear time in the number of elements of M .

2.5 Cut Polytope and Matroid Cycle Polytope

Given a graph G with edge set E , its *cut polytope* $CUT(G) \subseteq \mathbb{R}^E$ is the convex hull of the characteristic vectors of the cuts of G . For general graphs, a linear description of $CUT(G)$ is not known. However, for graphs without K_5 as a minor, $CUT(G)$ has been described by [6] as follows:

$$CUT(G) = \{x \in [0, 1]^E : x(F) - x(C \setminus F) \leq |F| - 1 \ \forall F \in \mathcal{F}\}, \quad (2.8)$$

where $\mathcal{F} = \{F \subset V(G) : F \subset C, C \text{ induced cycle of } G, |F| \text{ odd}\}$.

For a matroid $M = (E, \mathcal{B})$, a set $C \subseteq E$ is a *cycle* if $C = \emptyset$ or C is a disjoint union of circuits. The *cycle polytope* $C(M)$ of M is the convex hull of the characteristic vectors of its cycles [5]. Cycle polytopes can be seen as a generalization of cut polytopes. Indeed, it can be shown that if M is cographic, i.e. it is the dual of the forest matroid of some graph G , then the cycles of M correspond to the cuts of G , hence $C(M) = CUT(G)$.

A matroid is called binary if it can be represented over the finite field GF_2 . Given a matroid M , we denote by M^* its dual matroid. M is binary if and only if M^* is binary. An element e of a matroid is a chord of a circuit C if C is the symmetric difference of two circuits whose intersection is e . A chordless circuit is a circuit with no chords and the same definition can be applied to cocircuits, that are circuits in the dual matroid. F_7^* denotes the dual of the Fano matroid; R_{10} is a binary matroid associated with the 5×10 matrix whose columns are the 10 0/1 vectors with 3 ones and 2 zeros; $M_{K_5}^*$ is the dual of the forest matroid of K_5 .

In this section we prove Conjecture 2.1 for the cycle polytope $C(M)$ of the binary matroids M that have no minor isomorphic to F_7^* , R_{10} , $M_{K_5}^*$ and are 2-level. When those minors are forbidden, a complete linear description of the associated polytope is known (see [5]). This class includes all cut polytopes that are 2-level, and has been characterized in [44]:

Theorem 2.36. *Let M be a binary matroid with no minor isomorphic to F_7^* , R_{10} , $M_{K_5}^*$. Then $C(M)$ is 2-level if and only if M has no chordless cocircuit of length at least 5.*

Corollary 2.37. *The polytope $CUT(G)$ is 2-level if and only if G has no minor isomorphic to K_5 and no induced cycle of length at least 5.*

Recall that the cycle space of graph G is the set of its Eulerian subgraphs (subgraphs where all vertices have even degree), and it is known (see for instance [49]) to have a vector space structure over the field \mathbb{Z}_2 . This statement and one of its proofs easily generalizes to the cycle space (the set of all cycles) of binary matroids. We report the proof for completeness.

Lemma 2.38. *Let M be a binary matroid with d elements and rank r . Then the cycles of M form a vector space \mathcal{C} over \mathbb{Z}_2 with the operation of symmetric difference as sum. Moreover, \mathcal{C} has dimension $d - r$.*

Proof. That \mathcal{C} is a vector space can be easily verified using the fact that \mathcal{C} is closed under taking symmetric difference. This immediately derives from a characterization of binary matroids that can be found in [80], Theorem 9.1.2: M is binary if and only if the symmetric difference of any set of circuits is a disjoint union of circuits. We will now give a basis for \mathcal{C} of size $d - r$. The construction is analogous to the construction of a fundamental cycle basis in the cycle space of a graph. Consider a basis B of M . For any $e \in E \setminus B$, let C_e denote the unique circuit contained in $B + e$ (note that $e \in C_e$). Since $|B| = r$, we have a family $B_{\mathcal{C}} = \{C_{e_1}, \dots, C_{e_{d-r}}\}$ of the desired size. Note that the C_{e_i} 's are all linearly independent: indeed, C_e cannot be expressed as symmetric difference of other members of $B_{\mathcal{C}}$ since it is the only one containing e . We are left to show that $B_{\mathcal{C}}$ generates \mathcal{C} . Let $C \in \mathcal{C}$, $C \neq \emptyset$, and let $\{e_1, \dots, e_{d-r}\} \cap C = e_{i_1}, \dots, e_{i_k}$ for some $k \geq 1$ (indeed, $C \not\subseteq B$). Consider now $D = C \Delta C_{e_{i_1}} \Delta \dots \Delta C_{e_{i_k}}$. D is a cycle, however one can see that it is contained in B : for each $e \in E \setminus B$, if $e \in C$ then e appears exactly twice in the expression of D , hence $e \notin D$; if $e \notin C$, e does not appear in the expression at all. This implies that $D = \emptyset$, which is equivalent to $C = C_{e_{i_1}} \Delta \dots \Delta C_{e_{i_k}}$. \square

Chapter 2. On vertices and facets of 2-level polytopes arising in combinatorial settings

Corollary 2.39. *Let M be a binary matroid with d elements and rank r . Then M has exactly 2^{d-r} cycles.*

The only missing ingredient is a description of the facets of the cycle polytope for the class of our interest, which extends the description of the cut polytope given in (2.8).

Theorem 2.40. [5] *Let M be a binary matroid, and let $\overline{\mathcal{C}}$ be its family of chordless cocircuits. Then M has no minor isomorphic to F_7^* , R_{10} , $M_{K_5}^*$ if and only if*

$$C(M) = \{x \in [0, 1]^E : x(F) - x(C \setminus F) \leq |F| - 1 \text{ for } C \in \overline{\mathcal{C}}, F \subseteq C, |F| \text{ odd}\}.$$

Theorem 2.41. *Let M be a binary matroid with no minor isomorphic to F_7^* , R_{10} , $M_{K_5}^*$ and such that $C(M)$ is 2-level. Then $C(M)$ satisfies Conjecture 2.1.*

Proof. As remarked in [5] and [44], the following equations are valid for $C(M)$: a) $x_e = 0$, for e coloop of M ; and b) $x_e - x_f = 0$, for $\{e, f\}$ cocircuit of M .

The first equation is due to the fact that a coloop cannot be contained in a cycle, and the second to the fact that circuits and cocircuits have even intersection in binary matroids. A consequence of this is that we can delete all coloops and contract e for any cocircuit $\{e, f\}$ without changing the cycle polytope: for simplicity we will just assume that M has no coloops and no cocircuit of length 2. In this case $C(M)$ has full dimension $d = |E|$. Let r be the rank of M . Corollary 2.39 implies that $C(M)$ has 2^{d-r} vertices. Let now T be the number of cotriangles (i.e., cocircuits of length 3) in M , and S the number of cocircuits of length 4 in M . Thanks to Theorem 2.40 and to the fact that M has no chordless cocircuit of length at least 5, we have that $C(M)$ has at most $2d + 4T + 8S$ facets. Hence the bound we need to show is:

$$2^{d-r}(2d + 4T + 8S) \leq d2^{d+1} \iff 2T + 4S \leq d(2^r - 1).$$

Since the cocircuits of M are circuits in the binary matroid M^* , whose rank is $d - r$, we can apply Corollary 2.39 to get $T + S \leq 2^r - 1$, where the -1 comes from the fact that we do not count the empty set. Hence, if $d \geq 4$,

$$2T + 4S \leq 4(T + S) \leq d(2^r - 1).$$

The bound is loose for $d \geq 5$. The cases with $d \leq 4$ can be easily verified, the only tight examples being affinely isomorphic to cubes and cross-polytopes. \square

Corollary 2.42. *2-level cut polytopes satisfy Conjecture 2.1.*

2.6 On possible generalizations of the conjecture

So far, we provided a thorough analysis of 2-level polytopes coming from combinatorial settings. We hope that the reader shares with us the opinion that those polytopes are relevant for the mathematical community, and the 2-levelness property seem to be strong enough to leave hope for deep theorems on their structure. While we proved Conjecture 2.1 for all 2-level polytopes we could characterize, it remains open for the general case. Whether some techniques and ideas introduced in this paper can be extended to attack it also remains open. Here, we would like to discuss a different issue stemming from Conjecture 2.1: is 2-levelness the “right” assumption for proving $f_{d-1}(P)f_0(P) \leq d2^{d+1}$, or is this bound valid for a much more general class of 0/1 polytopes – or, more generally, of mathematical objects? We start the investigation of this question by providing some examples of “well-behaved” 0/1 polytopes that do not verify Conjecture 2.1. They can be seen as immediate generalizations of polytopes for which Conjecture 2.1 holds, see Corollary 2.27 and Proposition 2.10.

2.6.1 Forest polytope of $K_{2,n}$

Let P be the forest polytope of $K_{2,n}$. Note that P has dimension $d = 2n$. Conjecture 2.1 implies an upper bound of $n2^{2(n+1)} = O(4 + \varepsilon)^n$ for $f_0(P)f_{d-1}(P)$, for any $\varepsilon > 0$. Each subgraph of $K_{2,n}$ that takes, for each node v of degree 2, at most one edge incident to v , is a forest. Those graphs are 3^n . Moreover, each induced subgraph of $K_{2,n}$ that takes the nodes of degree n plus at least 2 other nodes is 2-connected, hence it induces a (distinct) facet of P . Those are $2^n - (n + 1)$. In total $f_0(P)f_{d-1}(P) = \Omega(6^n)$.

2.6.2 Spanning tree polytope of the skeleton of the 4-dimensional cube

Let G be the skeleton of the 4-dimensional cube, and P the associated spanning tree polytope. Through extensive computation⁵, we verified that $f_0(P)f_{d-1}(P) \geq 1.603 \cdot 10^{11}$, while the upper bound from Conjecture 2.1 is $\approx 1.331 \cdot 10^{11}$.

2.6.3 3-level min up/down polytopes

Fix $d \geq 3$. A 0/1 vector $x \in \{0, 1\}^d$ is “bad” if there are indices $0 < i < j < d$ such that $x_i = x_{j+1} = 1$ and $x_{i+1} = x_j = 0$. In other words, when seen as a bit-string, x is bad if it contains two or more separate blocks of 1’s. Let $P \subset \mathbb{R}^d$ be the convex hull of all 0/1 vectors that are not bad: this is a min up/down polytope, as defined in [72], with parameters $\ell_1 = 1$ and $\ell_2 = d - 1$ ⁶.

⁵We computed the number of spanning trees of G using the well known Kirchhoff’s matrix tree theorem [13]. The facets of the spanning tree polytope of a 2-connected graph G are roughly as many as the 2-connected, induced subgraphs of G whose contraction is 2-connected, and we compute them by exhaustive search. The Matlab code can be found at: <http://disopt.epfl.ch/files/content/sites/disopt/files/users/249959/facets.zip>

⁶Recall that the min up/down polytope is 2-level precisely when its parameters ℓ_1 and ℓ_2 are equal, and in that case the polytope satisfies Conjecture 2.1, see Proposition 2.10.

Chapter 2. On vertices and facets of 2-level polytopes arising in combinatorial settings

Each non-zero vertex x in P contains exactly one block of 1's, thus it is uniquely described by two indices $0 \leq i < j \leq d$, such that $x_k = 1$ if $i < k \leq j$, and $x_k = 0$ otherwise. Therefore (counting also the zero vector), P contains $\binom{d+1}{2} + 1$ vertices. On the other hand, from the facet characterization presented in [72] we know that

$$P = \left\{ x \in \mathbb{R}_+^d : \sum_{j=1}^k (-1)^{j-1} x_{i_j} \leq 1, \text{ for } 1 \leq i_1 < \dots < i_k \leq d \text{ s.t. } k \text{ is odd} \right\},$$

where all inequalities above are facet-defining. Moreover, since the polytope is full-dimensional (it contains the d -dimensional standard simplex) and no inequality is a multiple of another, they all define distinct facets. This means that there are d facets coming from non-negativity constraints, and 2^{d-1} facets that are in one-to-one correspondence with odd subsets of the index set $[d]$. Hence, the total number of facets is $2^{d-1} + d$. It is easy to check that for $d \geq 3$ we have

$$f_0(P) f_{d-1}(P) = \left[\binom{d+1}{2} + 1 \right] \cdot [2^{d-1} + d] > d 2^{d+1},$$

thus the polytope does not satisfy Conjecture 2.1. Note that P is a *3-level* polytope: for each facet F of P , there exist two translates of the affine hull of F such that all the vertices of P lie either in F or in one of those two translates.

In the remaining sections, we move to extensions of Conjecture 2.1 to other settings. In some of those cases we could prove that the conjecture does not hold. Others are interesting open questions.

2.6.4 Polytopes of minimum PSD rank

2-level polytopes are an example of polytopes with minimum PSD rank, i.e. such that they admit a semidefinite extension of size $1 + \dim(P)$, see [47]. A necessary and sufficient condition characterizing those polytopes is given in [47], where the full list of combinatorial classes of polytopes with minimum PSD rank in dimension 2 and 3 is also given. All those are combinatorially equivalent to some 2-level polytope of the same dimension, with the exception of the bipyramid over a triangle, which clearly verifies Conjecture 2.1. In [46], the full list of combinatorial classes of polytopes with minimum PSD rank in dimension 4 is given. By going through the list of their f -vectors in [46, Table 1], one easily checks that they also verify Conjecture 2.1. We are not aware of studies on higher-dimensional polytopes with minimum PSD rank. We remark that in [48] it is proved that matroid base polytopes have minimum PSD rank if and only if they are 2-level, hence Theorem 2.26 trivially implies that all matroid polytopes with minimum PSD rank satisfy the bound of Conjecture 2.1.

2.6.5 Polytopes with structured linear relaxations

We now consider another possible generalization of the conjecture, based on the \mathcal{H} -embedding of 2-level polytopes as defined in [10]. It is shown in [10] that the family of 2-level polytopes of dimension d is affinely equivalent to the family of integral polytopes of the form

$$P = \{x \in \mathbb{R}^d : 0 \leq x(I) \leq 1 \quad \text{for all } I \subseteq \mathcal{I}\} \quad (2.9)$$

for some $\mathcal{I} \subseteq 2^{[d]}$. Hence, Conjecture 2.1 holds for 2-level polytopes if and only if it holds for integral polytopes of the form (2.9). First, notice that the bound of the conjecture does not hold for general (i.e. not integral) polytopes of the form (2.9), as for instance fractional stable set polytopes are of such form. In particular, consider the fractional stable set polytope of the complete graph on d vertices, $P = \{x \in \mathbb{R}^d : x \geq 0, x_i + x_j \leq 1 \ \forall i \neq j, i, j \in [d]\}$. Clearly P can be written in form (2.9), and has $d + \binom{d}{2}$ facets. It is not hard to see (we refer to [90, Chapter 64] for further details) that P has $d + 1$ integral vertices, and exponentially many fractional vertices obtained by setting at least three coordinates to $1/2$ and the others to 0, hence $f_0(P) = 2^d - \binom{d}{2}$, and we have $f_0(P)f_{d-1}(P) = \left[d + \binom{d}{2}\right] \left[2^d - \binom{d}{2}\right] > d2^{d+1}$ for $d \geq 5$.

Another natural question is whether the bound of Conjecture 2.1 holds for integral polytopes that admit a *linear relaxation* of the kind (2.9). More formally, let P_I be the integer hull of a polytope P . Is it true that, for all P of the form (2.9), one has $f_0(P_I)f_{d-1}(P_I) \leq d2^{d+1}$? Note that this seems to be too general to be true, since such P include, for instance, all stable set polytopes. However, given the difficulty of building explicit polytopes with many facets (see [68] for some constructions and a discussion), finding a counter-example is non-trivial. Through extensive computation with `polymake`, we found a 12-dimensional polytope P that violates the conjecture. Indeed, for $d = 12$ the bound of the conjecture is 98304, while $f_0(P_I)f_{d-1}(P_I) = 535392$. We give an explicit description of the polytope in the appendix.

2.6.6 0/1 matrices generalizing slack matrices of 2-level polytopes

As mentioned in Section 2.1, Conjecture 2.1 can be rephrased as an upper bound on the number of entries of the smallest slack matrices of 2-level polytopes. It is then a natural question whether one can extend the conjecture on classes of matrices strictly containing those matrices.

Let $M \in \{0, 1\}^{m \times n}$ be a matrix without any repeated row or column. Using the characterization given in [43], we have that M is the slack matrix of a 2-level d -polytope P with $d \geq 2$ if and only if:

- (i) $rk(M) = d + 1$;
- (ii) The vector with all components equal to 1 belongs to the space generated by the rows of M ;

- (iii) The cone generated by the rows of M coincide with the intersection of the space generated by the rows of M with the nonnegative orthant.

Moreover, if M is a minimal slack matrix for P , then (iv) rows of M have incomparable supports and (v) columns of M have incomparable supports. We want to understand what happens to Conjecture 2.1 when one of those properties is relaxed.

Relaxing (i) does not make sense, since it leads to slack matrices of 2-level polytopes of any dimension, which clearly violate the conjecture. Now suppose we relax (iv). Let A be the slack matrix of the d -dimensional cube, and A' obtained from A by adding a row of 1s. A' verifies properties (i)-(ii)-(iii)-(v), since it is obtained from A by adding a row that is already in the conic hull of rows of A . On the other hand, since the cube verifies Conjecture (2.1) at equality, A' does not verify the conjecture. Similarly, if we relax (v) instead of (iv), add a column of 1 to A as to obtain A'' . Note that this new column is also in the conic hull of the columns of A , since A^T is the slack matrix of the d -dimensional cross-polytope. Hence A'' verifies properties (i)-(ii)-(iii)-(iv) but not Conjecture 2.1. Finding counterexamples to the conjecture when property (ii) or (iii) are relaxed seems to be harder, hence an interesting open question. Note that, when (ii) is relaxed, A is the slack matrix of a polyhedral cone, see again [43].

We now investigate what happens if we relax the conditions above even further, and only impose that the rank of $M \in \{0, 1\}^{m \times n}$ is d , and that M does not have any repeated row or column. From the discussion above, we know that M does not verify Conjecture 2.1, but which bound can one give on $m \cdot n$? A standard argument implies that the maximum number of distinct rows (resp. columns) is 2^d , hence $m \cdot n \leq 4^d$. Indeed, consider c_1, \dots, c_d linearly independent columns of M . Any other column is a linear combinations of the c_i 's. But then, if two rows coincide on c_1, \dots, c_d , then they are equal, a contradiction. Hence all the rows must be distinct on c_1, \dots, c_d , but then, being M 0/1, there can be at most 2^d rows.

We now show that the bound $m \cdot n \leq 4^d$ is not tight. We first show the following:

Claim 2.43. Let M be a 0/1 matrix of rank d , containing the identity matrix I_d as submatrix, and with no repeated rows and columns; then M has size at most $(d + 1)2^d$.

Proof. Assume for simplicity (and without loss of generality) that I_d is contained in the upper left corner of M . Then the first d columns of M , denoted by c_1, \dots, c_d , are linearly independent, and the first d entries of c_i form the vector e_i for $i = 1, \dots, d$. Since $\text{rk}(M) = d$, every other column c_i , $i > d$, can be written as $\sum_{j=1}^d \alpha_j^{(i)} c_j$ for some coefficients $\alpha_j^{(i)}$'s. But the first d entries of such c_i are exactly $\alpha_1^{(i)}, \dots, \alpha_d^{(i)}$, hence, as M is 0/1, we have $\alpha_j^{(i)} \in \{0, 1\}$ for any i, j . Now, consider a graph G with vertex set $[d]$, where node j and node k are adjacent if, for some i , we have $\alpha_j^{(i)} = \alpha_k^{(i)} = 1$. Clearly each column of M corresponds to a clique of G (including c_1, \dots, c_d , which correspond to singletons). Notice also that two columns c_i, c_h cannot correspond to the same clique, as this would imply that $\alpha^{(i)} = \alpha^{(h)}$, hence that $c_i = c_h$. Now, for any row r of M , consider its first d entries. If for some $j < k \leq d$ we have $r_j = r_k = 1$, then we cannot

2.6. On possible generalizations of the conjecture

have $\alpha_j^{(i)} = \alpha_k^{(i)} = 1$ for any column i , otherwise the entry of M corresponding to r and c_i would be at least 2, a contradiction. Hence, each row of M corresponds to a stable set of G . As before, notice that no two rows can correspond to the same stable set. But then we can apply Corollary 2.6 to G : defining $\mathcal{C}', \mathcal{S}'$ for G as in the corollary, we obtain that the size of M is at most $|\mathcal{C}'||\mathcal{S}'| \leq (d+1)2^d$. \square

Now, it easily follows that any 0/1 matrix with rank d , and no repeated rows or columns, cannot have 2^d rows and 2^d columns. Assume that M has 2^d rows: we will show that it satisfies the hypothesis of the above claim, i.e. that it contains I_d as a submatrix. Let c_1, \dots, c_d linearly independent columns of M , and let M' be M restricted to these columns. As argued before, the rows of M' are all different: two rows that coincide in M' yield equal rows in M . But then all possible 0/1 vectors must appear as rows of M' , in particular M' (hence M) contains I_d as a submatrix. In conclusion, the claim implies that M has at most $d+1$ columns, and analogously, if we assume that M has 2^d columns, we obtain that M has at most $d+1$ rows, hence the bound 4^d cannot be tight.

One might wonder whether we can apply the above claim to the slack matrix of some interesting 2-level polytopes, to bound its size. However, we now show that the hypotheses of the claim are too strong to be satisfied by any interesting slack matrix. Let M be a minimal 0/1 slack matrix of a polytope P of dimension d , hence $\text{rk}(M) = d+1$, and assume that M contains I_{d+1} . We claim that P is the $d+1$ -dimensional simplex and $M = I_{d+1}$. The argument is similar to the previous one and we only sketch it. Condition (ii) states that the vector with all components equal to 1 belongs to the space generated by the rows of M . But this space is generated by those rows r_1, \dots, r_{d+1} of M which contain I_{d+1} , hence $\vec{1} = \sum_{i=1}^{d+1} \alpha_i r_i$, which implies similarly as before that $\alpha_i = 1$ for $i = 1, \dots, d+1$. It then follows from the fact that M has all distinct columns that M has exactly $d+1$ columns. Hence P is a d -dimensional polytope with $d+1$ vertices, i.e. it is a simplex.

3 On the extension complexity of the stable set polytope of bipartite graphs

3.1 Introduction

We recall from Chapter 1 that the extension complexity $\text{xc}(P)$ of P is the minimum number of facets of any extension of P . If Q is an extension of P with significantly fewer facets than P , then it is advantageous to run linear programming algorithms over Q instead of P .

One example of a polytope that admits a much more compact representation in a higher dimensional space is the spanning tree polytope, $\text{STP}(G)$. Edmonds' [26] classic description of $\text{STP}(G)$ has $2^{\Omega(|V|)}$ facets. However, Wong [99] and Martin [79] proved that for every connected graph $G = (V, E)$,

$$|E| \leq \text{xc}(\text{STP}(G)) \leq O(|V| \cdot |E|).$$

Fiorini, Massar, Pokutta, Tiwary, and de Wolf [36] were the first to show that many polytopes arising from NP-hard problems (such as the stable set polytope) do indeed have high extension complexity. Their results answer an old question of Yannakakis [100] and do not rely on any complexity assumptions such as $P \neq NP$.

On the other hand, Rothvoß [86] proved that the perfect matching polytope of the complete graph K_n has extension complexity at least $2^{\Omega(n)}$. This is somewhat surprising since the maximum weight matching problem can be solved in polynomial-time via Edmond's blossom algorithm [27]. By now many accessible introductions to extension complexity are available (see [59], [20], [21], [87]).

Let $G = (V, E)$ be a (finite, simple) graph with $n := |V|$ and $m := |E|$, and let $\text{STAB}(G)$ be its stable set polytope, as defined in Section 1. As previously mentioned, $\text{STAB}(G)$ can have very high extension complexity. In [36], it is proved that if G is obtained from a complete graph by subdividing each edge twice, then $\text{xc}(\text{STAB}(G))$ is at least $2^{\Omega(\sqrt{n})}$. Recently, Göös, Jain, and Watson [41] improved this to $2^{\Omega(n/\log n)}$, via a different class of graphs. For perfect graphs, Yannakakis [100] proved an upper bound of $n^{O(\log n)}$, and it is an open problem whether

Yannakakis' upper bound can be improved to a polynomial bound.

In what follows we restrict our attention to bipartite graphs, which are perfect. Let $G = (V, E)$ be a bipartite graph with n vertices, m edges and no isolated vertices. By total unimodularity,

$$\text{STAB}(G) = \{x \in \mathbb{R}^V \mid x_u \geq 0 \text{ for all } u \in V, x_u + x_v \leq 1 \text{ for all } uv \in E\},$$

and so $n \leq \text{xc}(\text{STAB}(G)) \leq n + m$. In this case $\text{xc}(\text{STAB}(G))$ lies in a very narrow range, and it is a good test of current methods to see if we can improve these bounds.

The situation is analogous to what happens with the spanning tree polytope of (arbitrary) graphs, where as previously mentioned, we also know that $\text{xc}(\text{STP}(G))$ lies in a very narrow range. Indeed, a notorious problem of Goemans (see [64]) is to improve the known bounds for $\text{xc}(\text{STP}(G))$, but this is still wide open. However, for the stable set polytopes of bipartite graphs, we are able to improve on the known bounds, as we summarize below.

Contribution and organization. The main results of the chapter are the following.

- We improve over the trivial upper bound on $\text{xc}(\text{STAB}(G))$, G bipartite: in particular we prove that for any bipartite graph G with n vertices, the extension complexity of $\text{STAB}(G)$ is $O(n^2 / \log n)$. This is an improvement when G has quadratically many edges.
- We improve over the trivial lower bound on $\text{xc}(\text{STAB}(G))$, G bipartite: in particular we find an infinite class of graphs such that the stable set polytope of every n -vertex graph in the class has extension complexity $\Omega(n \log n)$. These are the first known examples of stable set polytopes of bipartite graphs where the extension complexity is more than linear in the number of vertices. For instance, $\text{xc}(\text{STAB}(K_{n/2, n/2})) = \Theta(n)$. As argued in Chapter 1, our lower bound (see Theorem 3.11) is significant because no other non-trivial lower bound on extension complexity is known even for the more general classes of stable set polytopes of perfect graphs and of 2-level polytopes.
- We show that it is not possible to prove a better lower bound for our class of graphs using our approach. At the core of this result lies a combinatorial problem which has been studied independently by Kaibel and Averkov ([60]) under the name of ‘class covering’, in the context of rectangle covers for the spanning tree polytope. We briefly describe this connection and prove a lower bound on the size of class covers.

In Section 3.2 we define rectangle covers and fooling sets and we give examples of 3-regular graphs with tight fooling sets. We prove Theorem 3.5 in Section 3.3 and Theorem 3.11 in Section 3.4. In Section 3.5 we show that it is impossible to prove a better lower bound with the approach in Section 3.4. Thus, to further improve the lower bound, different methods (or different graphs) are required. We conclude the Chapter with Section 3.6, where we describe the connection between a covering problem described in Section 3.5 and the class cover problem.

3.2 Rectangle Covers and Fooling Sets

Consider a polytope $P := \text{conv}(X) = \{x \in \mathbb{R}^d \mid Ax \geq b\}$, where $X := \{x^{(1)}, \dots, x^{(n)}\} \subseteq \mathbb{R}^d$, $A \in \mathbb{R}^{m \times d}$ and $b \in \mathbb{R}^m$. Recall from Section 1 that the slack matrix of P (with respect to the chosen inner and outer descriptions of the polytope) is the matrix $S \in \mathbb{R}_{\geq 0}^{m \times n}$ having rows indexed by the inequalities $A_1 x \geq b_1, \dots, A_m x \geq b_m$ and columns indexed by the points $x^{(1)}, \dots, x^{(n)}$, defined as $S_{ij} := A_i x^{(j)} - b_i \geq 0$.

We recall that Yannakakis [100] proved that the extension complexity of P equals the nonnegative rank of S (see Theorem 1.5). In this chapter, we only rely on a lower bound that follows directly from this fact. For a matrix M , we define the *support* of M as $\text{supp}(M) := \{(i, j) \mid M_{ij} \neq 0\}$. A *rectangle* is any set of the form $R = I \times J$, with $R \subseteq \text{supp}(M)$. A *size- k rectangle cover* of M is a collection R_1, \dots, R_k of rectangles such that $\text{supp}(M) = R_1 \cup \dots \cup R_k$. The *rectangle covering bound* of M is the minimum size of a rectangle cover of M , and is denoted $\text{rc}(M)$.

Theorem 3.1 (Yannakakis, [100]). *Let P be a polytope with $\dim(P) \geq 1$ and let S be any slack matrix of P . Then, $\text{xc}(P) \geq \text{rc}(S)$.*

A *fooling set* for M is a set of entries $F \subseteq \text{supp}(M)$ such that $M_{i\ell} \cdot M_{kj} = 0$ for all distinct $(i, j), (k, \ell) \in F$. The largest size of a fooling set of M is denoted by $\text{fool}(M)$. Clearly, $\text{rc}(M) \geq \text{fool}(M)$.

Let G be a bipartite graph. The *edge vs stable set matrix* of G , denoted $M(G)$, is the 0/1 matrix with a row for each edge of G , a column for each stable set of G , and a 1 in position (e, S) if and only if $e \cap S = \emptyset$ (as usual, we regard edges as pairs of vertices). We say that G has a *tight fooling set* if $M(G)$ has a fooling set of size $|E(G)|$. Note that if G has a tight fooling set, then the non-negative rank of $M(G)$ is exactly $|E(G)|$. Also observe that the property of having a tight fooling set is closed under taking subgraphs.

It is easy to check that even cycles have tight fooling sets. We now give an infinite family of 3-regular graphs that have tight fooling sets. A graph is C_4 -free if it does not contain a cycle of length four.

Theorem 3.2. *Let $G = (V, E)$ be a 3-regular, C_4 -free bipartite graph. Then G has a tight fooling set.*

Proof. For $X \subseteq V$, we let $N(X)$ denote the set of neighbours of X . Let $V = A \cup B$ be a bipartition of the vertex set, and let $\phi : E \rightarrow \{1, 2, 3\}$ be a proper edge coloring of G , which exists by 3-regularity and König's edge-coloring theorem (see e.g. [91, Theorem 20.1]). For each vertex $a \in A$, we name its neighbors $a_1, a_2, a_3 \in B$ so that $\phi(aa_i) = i$. For each $a \in A$, consider the

following stable sets:

$$\begin{aligned} S_{aa_1} &:= A \setminus \{a\} \\ S_{aa_2} &:= \{a_1\} \cup \{a' \in A \mid a' \notin N(a_1)\} \\ S_{aa_3} &:= B \setminus \{a_3\}. \end{aligned}$$

This defines a stable set S_e disjoint from e , for every edge $e \in E$. Since ϕ is proper, no two of these stable sets are equal. We claim that $\{(e, S_e) \mid e \in E\}$ is a fooling set in the edge vs stable set matrix of G .

Let e and f be distinct edges. We want to show that S_e intersects f or S_f intersects e . Consider the following three cases. Let $e = aa_i$, where $i = \phi(e)$.

Case 1. If $\phi(e) = 1$, then $S_e = S_{aa_1}$ intersects f unless $f = aa_i$ for some $i \in \{2, 3\}$. In both cases we have $a_1 \in S_f \cap e$.

Case 2. If $\phi(e) = 3$, then $S_e = S_{aa_3}$ intersects f unless $f = a'a_3$ for some $a' \in A$. Either $\phi(f) = 1$ and S_f intersects e (as in Case 1), or $\phi(f) = 2$. In the last case, since G is C_4 -free, we have $a \notin N(a'_1)$. It follows that $S_f = S_{a'a_3} = S_{a'a'_2}$ intersects e .

Case 3. If $\phi(e) = 2$, then we may also assume $\phi(f) = 2$ since otherwise by exchanging the roles of e and f we are back to one of the previous cases. Let a' denote the endpoint of f in A , so that $f = a'a'_2$. Because ϕ is proper, $a' \neq a$ and $a'_1 \neq a_1$. Since G is C_4 -free, we have $a \notin N(a'_1)$ or $a' \notin N(a_1)$. Hence, $a \in S_f \cap e$ or $a' \in S_e \cap f$. \square

Note that there are infinitely many 3-regular, C_4 -free bipartite graphs. For example, we can take a hexagonal grid on a torus.

3.3 An Improved Upper Bound

In this section we prove Theorem 3.5. We use the following result of Martin [79].

Lemma 3.3. *If Q is a nonempty polyhedron, $\gamma \in \mathbb{R}$, and*

$$P = \{x \mid \langle x, y \rangle \leq \gamma \text{ for every } y \in Q\},$$

then $\text{xc}(P) \leq \text{xc}(Q) + 1$.

The *edge polytope* $P_{\text{edge}}(G)$ of a graph G is the convex hull of the incidence vectors in $\mathbb{R}^{V(G)}$ of all edges of G . The second ingredient we need is the following bound on the extension complexity of the edge polytope of all n -vertex graphs due to Fiorini, Kaibel, Pashkovich, and Theis [33, Lemma 3.4]. This bound follows from a nice result of Tuza [96], which states that

every n -vertex graph can be covered with a set of bicliques of total weight $O(n^2 / \log n)$, where the weight of a biclique is its number of vertices.

Lemma 3.4. *For every graph G with n vertices, $\text{xc}(\text{P}_{\text{edge}}(G)) = O(n^2 / \log n)$.*

We are now in position to prove our upper bound:

Theorem 3.5. *For all bipartite graphs G with n vertices, the extension complexity of $\text{STAB}(G)$ is $O(n^2 / \log n)$.*

Proof. Let $G = (V, E)$. Since

$$\text{STAB}(G) = \mathbb{R}_{\geq 0}^V \cap \{x \in \mathbb{R}^V \mid \langle x, y \rangle \leq 1 \text{ for every } y \in \text{P}_{\text{edge}}(G)\},$$

By Lemmas 3.3 and 3.4, the extension complexity of $\text{STAB}(G)$ is $O(n^2 / \log n)$. \square

3.4 An Improved Lower Bound

In this section we describe a class of bipartite graphs whose stable set polytope has super-linear extension complexity. The examples we use are incidence graphs of finite projective planes. We will not use any theorems from projective geometry, but the interested reader can refer to [25].

Let q be a prime power, $\text{GF}(q)$ be the field with q elements, and $\text{PG}(2, q)$ be the projective plane over $\text{GF}(q)$. The *incidence graph* of $\text{PG}(2, q)$, denoted $\mathcal{I}(q)$, is the bipartite graph with bipartition $(\mathcal{P}, \mathcal{L})$, where \mathcal{P} is the set of points of $\text{PG}(2, q)$, \mathcal{L} is the set of lines of $\text{PG}(2, q)$, and $p \in \mathcal{P}$ is adjacent to $\ell \in \mathcal{L}$ if and only if the point p lies on the line ℓ . For example, $\text{PG}(2, 2)$ and its incidence graph $\mathcal{I}(2)$ are depicted in Figure 3.1.




Figure 3.1 – $\text{PG}(2, 2)$ and its incidence graph $\mathcal{I}(2)$.

Before proving our lower bound we gather a few lemmas on binomial coefficients. The first two are well-known, so we omit the easy proofs.

Lemma 3.6. *For all integers h and c with $h \geq c \geq 0$*

$$\sum_{j=c}^h \binom{j}{c} = \binom{h+1}{c+1}.$$

Lemma 3.7. *For all positive integers x, y , and h ,*

$$\sum_{j=0}^h \binom{x+j}{j} \binom{h+y-j}{h-j} = \binom{x+y+h+1}{h}.$$

Lemma 3.8. *Let q, c, t be positive integers with $c+t \leq q+1$. Then*

$$t \sum_{k=c}^{q+1-t} \frac{1}{k} \binom{q+1-t-c}{k-c} \binom{q}{k}^{-1} = \binom{t+c-1}{t}^{-1} \leq \frac{1}{c}.$$

Proof. We have that

$$\begin{aligned} & t \sum_{k=c}^{q+1-t} \frac{1}{k} \binom{q+1-t-c}{k-c} \binom{q}{k}^{-1} \\ &= \frac{t(q+1-t-c)!}{q!} \sum_{k=c}^{q+1-t} \frac{(k-1)!(q-k)!}{(k-c)!(q+1-t-k)!} \\ &= \frac{t(q+1-t-c)!}{q!} (c-1)!(t-1)! \sum_{k=c}^{q+1-t} \binom{k-1}{c-1} \binom{q-k}{t-1}. \end{aligned}$$

Moreover,

$$\begin{aligned} \sum_{k=c}^{q+1-t} \binom{k-1}{c-1} \binom{q-k}{t-1} &= \sum_{j=0}^{q+1-t-c} \binom{c-1+j}{c-1} \binom{q-c-j}{t-1} \\ [h = q+1-t-c, x = c-1, y = t-1] &= \sum_{j=0}^h \binom{x+j}{j} \binom{h+y-j}{h-j} \\ [\text{by Lemma 3.7}] &= \binom{x+y+h+1}{h} \\ &= \binom{q}{q+1-t-c}. \end{aligned}$$

We conclude that

$$\begin{aligned} t \sum_{k=c}^{q+1-t} \frac{1}{k} \binom{q+1-t-c}{k-c} \binom{q}{k}^{-1} &= \frac{t(q+1-t-c)!}{q!} \frac{q!(c-1)!(t-1)!}{(q+1-t-c)!(t+c-1)!} \\ &= \binom{t+c-1}{t}^{-1}. \end{aligned}$$

The number of t -subsets of a set of size $t + c - 1$ is at least c , since it includes all t -subsets containing a fixed set of size $t - 1$. Hence, $\binom{t+c-1}{t}^{-1} \leq \frac{1}{c}$. \square

From the definition of $\text{PG}(2, q)$ it follows that that $\mathcal{J}(q)$ is $(q + 1)$ -regular, $|V(\mathcal{J}(q))| = 2(q^2 + q + 1)$, and $|E(\mathcal{J}(q))| = (q + 1)(q^2 + q + 1)$. Let $n = q^2 + q + 1$ and note that $\mathcal{J}(q)$ has $2n$ vertices. We let \mathcal{P} and \mathcal{L} denote the set of points and lines of $\text{PG}(2, q)$. We also use the fact that $\mathcal{J}(q)$ is C_4 -free.

We denote the edge vs stable set incidence matrix of $\mathcal{J}(q)$ by S_q . Each 1-entry of S_q is of the form (e, S) where $e \in E$, $S \subseteq V$ is a stable set, and $e \cap S = \emptyset$. To prove Theorem 3.11 we will assign weights to the 1-entries of S_q in such a way that the total weight is at least $\Omega(n \log n)$, while the weight of every rectangle is at most 1. The only entries that will receive non-zero weight are what we call *special entries*, which we now define.

Definition 3.9. A 1-entry of S_q is *special* if it has the form $(e, S(X))$ where

- $e = p\ell$ with $p \in \mathcal{P}, \ell \in \mathcal{L}$,
- $X \subseteq N(\ell) \setminus \{p\}$, X non-empty,
- $S(X) = X \cup (\mathcal{L} \setminus N(X))$.

We also need the following compact representation of maximal rectangles.

Definition 3.10. Let R be a maximal rectangle. Then R is determined by a pair $(\mathcal{P}_R, \mathcal{L}_R)$ with $\mathcal{P}_R \subseteq \mathcal{P}$, $\mathcal{L}_R \subseteq \mathcal{L}$, where the rows of R are all the edges between \mathcal{P}_R and \mathcal{L}_R and the columns of R are all the stable sets $S \subseteq V \setminus (\mathcal{P}_R \cup \mathcal{L}_R)$.

We are now ready to prove our lower bound.

Theorem 3.11. *Let q be a prime power and $n = q^2 + q + 1$. Then there exists a constant $c > 0$ such that*

$$\text{xc}(\text{STAB}(\mathcal{J}(q))) \geq cn \log n.$$

Proof. Let $n = q^2 + q + 1$. Let $V = \mathcal{P} \cup \mathcal{L}$ be the vertices of $\mathcal{J}(q)$, and E be the edges of $\mathcal{J}(q)$. To each special entry $(e, S(X))$ we assign the weight

$$w(e, S(X)) = \frac{1}{|X| \binom{q}{|X|} (q + 1)}.$$

All other entries of S_q receive weight zero.

Claim 3.12. $w(S_q) := \sum_{(e, S)} w(e, S) \geq cn \log n$ for some constant c .

Subproof. We have that

$$\begin{aligned} \sum_{(e,S)} w(e,S) &= \sum_{(e,S(X)) \text{ special}} w(e,S(X)) = \sum_{e \in E} \sum_{k=1}^q \binom{q}{k} \frac{1}{k \binom{q}{k} (q+1)} \\ &= \frac{|E|}{q+1} \sum_{k=1}^q \frac{1}{k} = n \sum_{k=1}^q \frac{1}{k} > cn \log n. \end{aligned}$$

The claim follows. ■

Let $R = (\mathcal{P}_R, \mathcal{L}_R)$ be an arbitrary maximal rectangle. We finish the proof by showing that $w(R) := \sum_{(e,S) \in R} w(e,S) \leq 1$. Together with Claim 3.12 this clearly implies Theorem 3.11. We will need the following obvious but useful Claim.

Claim 3.13. A special entry $(p\ell, S(X))$ is covered by $R = (\mathcal{P}_R, \mathcal{L}_R)$ if and only if $X \cap \mathcal{P}_R = \emptyset$, $\mathcal{L}_R \subseteq N(X)$, $p \in \mathcal{P}_R$, and $\ell \in \mathcal{L}_R$.

We consider two cases. First suppose that $\mathcal{L}_R = \{\ell\}$ for some ℓ . Then the only special entries covered by R are of the form $(p\ell, S(X))$, with $X \subseteq N(\ell) \setminus \mathcal{P}_R$. Let $N(\ell) \cap \mathcal{P}_R = \{p_1, \dots, p_t\}$, where $1 \leq t \leq q+1$. To compute $w(R)$ we have to sum over all edges $p_i\ell$ and over all subsets $X \subseteq N(\ell) \setminus \{p_1, \dots, p_t\}$. It follows that

$$\begin{aligned} w(R) &= \sum_{i=1}^t \sum_{k=1}^{q+1-t} \binom{q+1-t}{k} \frac{1}{k \binom{q}{k} (q+1)} \\ &= t \sum_{k=1}^{q+1-t} \frac{(q+1-t)!}{k!(q+1-t-k)!} \frac{k!(q-k)!}{kq!(q+1)} \\ &= \frac{t(q+1-t)!(t-1)!}{(q+1)!} \sum_{k=1}^{q+1-t} \binom{q-k}{q+1-t-k} \frac{1}{k} \\ &= \frac{1}{\binom{q+1}{t}} \sum_{k=1}^{q+1-t} \binom{q-k}{t-1} \frac{1}{k} \leq \frac{1}{\binom{q+1}{t}} \sum_{j=t-1}^{q-1} \binom{j}{t-1} = \frac{1}{\binom{q+1}{t}} \binom{q}{t} \leq 1, \end{aligned}$$

where the last equality follows from Lemma 3.6.

The remaining case is if $|\mathcal{L}_R| \geq 2$. For $\ell \in \mathcal{L}_R$ such that $(p\ell, S(X))$ is covered by R for some p, X , define

$$k_\ell = \min\{|X| \mid \text{there exist } p, X : (p\ell, S(X)) \text{ is a special entry covered by } R\}.$$

Claim 3.14. Let $(p\ell, S(X))$ be a special entry covered by R such that $|X| = k_\ell$. Then for each p', Y such that R covers $(p'\ell, S(Y))$, we have $X \subseteq Y$.

Subproof. For each $\ell' \in \mathcal{L}_R \setminus \{\ell\}$ (there is at least one since $|\mathcal{L}_R| > 1$), we have $\ell' \in N(X)$ by Claim 3.13. That is, there is $x = x(\ell') \in X$ adjacent to ℓ' . Similarly, since $\ell' \in N(Y)$, there is $y = y(\ell') \in Y$ adjacent to ℓ' . Now, if $x(\ell') \neq y(\ell')$, then $\mathcal{S}(q)$ contains a 4-cycle, which is a

contradiction. Hence we must have $x(\ell') = y(\ell')$ for all $\ell' \in \mathcal{L}_R \setminus \{\ell\}$. Now if there is an $x \in X$ such that $x \neq x(\ell')$ for every $\ell' \in \mathcal{L}_R \setminus \{\ell\}$, then $(p\ell, S(X \setminus \{x\}))$ is still covered by R , contradicting the minimality of X . We conclude $X \subseteq Y$, as required. ■

Now fix $\ell \in \mathcal{L}_R$, and let

$$w(\ell) = \sum \{w(p\ell, S(X)) \mid (p\ell, S(X)) \text{ special}\}.$$

Claim 3.15. For every $\ell \in \mathcal{L}_R$,

$$w(\ell) \leq \frac{1}{(q+1)k_\ell}.$$

Subproof. Let $N(\ell) \cap \mathcal{P}_R = \{p_1, \dots, p_t\}$, where $1 \leq t \leq q+1$. Let X be such that $(p\ell, S(X))$ is a special entry covered by R and $|X| = k_\ell$. By Claim 3.14, the only special entries appearing in the above sum are of the form $(p_i\ell, S(Y))$ where $i \in [t]$ and $X \subseteq Y \subseteq (\mathcal{P} \setminus \mathcal{P}_R) \cap N(\ell)$. Therefore

$$w(\ell) \leq t \sum_{k=k_\ell}^{q+1-t} \binom{q+1-t-k_\ell}{k-k_\ell} \frac{1}{k \binom{q}{k} (q+1)} \leq \frac{1}{(q+1)k_\ell},$$

where the last inequality follows from Lemma 3.8 with $c = k_\ell$. ■

Claim 3.16. For every $\ell \in \mathcal{L}_R$, $|\mathcal{L}_R| \leq (q+1)k_\ell$.

Subproof. Again, let X be such that $(p\ell, S(X))$ is covered by R and assume that $|X| = k_\ell$. By Claim 3.13, we have $\mathcal{L}_R \subseteq N(X)$.

Hence $|\mathcal{L}_R| \leq |N(X)| \leq (q+1)|X| = (q+1)k_\ell$. ■

By Claim 3.15 and Claim 3.16, for every $\ell \in \mathcal{L}_R$, $w(\ell) \leq \frac{1}{|\mathcal{L}_R|}$. But clearly $w(R) = \sum_{\ell \in \mathcal{L}_R} w(\ell)$, and so $w(R) \leq 1$, as required. This completes the entire proof. □

3.5 A small rectangle cover of the special entries

In this section we show that the submatrix of special entries considered in the previous section has a rectangle cover of size $O(n \log n)$. Combined with Theorem 3.11, this implies that a minimal set of rectangles that cover all the special entries always has size $\Theta(n \log n)$. Thus, to improve our bound, we must consider a different set of entries of the slack matrix, or use a different set of graphs.

This cover will be built from certain labeled trees which we now define. Note that the *length* of a path is its number of edges.

Definition 3.17. For every integer $k \geq 1$, we build a tree $T(k)$ recursively:

- The tree $T(1)$ consists of a root r and a single leaf attached to it.
- For $k > 1$, we construct $T(k)$ by first identifying one end of a path P_1 of length $k_1 := \left\lceil \frac{k}{2} \right\rceil$ to another end of a path P_2 of length $k_2 := \left\lfloor \frac{k}{2} \right\rfloor$ along a root vertex r . Let λ_i be the end of P_i that is not r . We then attach a copy of $T(k_i)$ to λ_{3-i} , identifying λ_{3-i} with the root of $T(k_i)$. We call P_1 and P_2 the *main paths* of $T(k)$.

The next Lemma follows easily by induction on k .

Lemma 3.18. *For all $k \geq 1$,*

1. $T(k)$ has $O(k \log k)$ vertices;
2. $T(k)$ has k leaves;
3. every path from the root r to a leaf has length k .

Definition 3.19. We recursively define a labeling $\varphi_k : V(T(k)) \setminus \{r\} \rightarrow [k]$ as follows:

- Let v be the non-root vertex of $V(T(1))$ and set $\varphi_1(v) := 1$.
- For $k > 1$, let P_1 and P_2 be the main paths of $T(k)$. We name the vertices of P_1 as $r, v_1, \dots, v_{\lceil \frac{k}{2} \rceil}$ and P_2 as $r, v_{\lceil \frac{k}{2} \rceil + 1}, \dots, v_k$, where these vertices are listed according to their order along P_1 and P_2 . Set $k_1 := \left\lceil \frac{k}{2} \right\rceil$ and $k_2 := \left\lfloor \frac{k}{2} \right\rfloor$. Note that $V(T(k)) = \bigcup_{i=1,2} (V(P_i) \cup V(B_i))$, where B_i is a copy of the tree $T(k_{3-i})$. We define

$$\varphi_k(v) = \begin{cases} i, & \text{if } v = v_i \\ \varphi_{k_2}(v) + k_1, & \text{if } v \in V(B_1) \setminus V(P_1) \\ \varphi_{k_1}(v), & \text{if } v \in V(B_2) \setminus V(P_2) \end{cases}$$

For each vertex $v \in T(k)$ we let $P(v)$ be the path in $T(k)$ from r to v .

Lemma 3.20. *Let φ_k , B_1 , and B_2 be as in Definition 3.19.*

1. If L is the set of leaves of $T(k)$, then $\varphi_k(L \cap V(B_1)) = \{\left\lceil \frac{k}{2} \right\rceil + 1, \dots, k\}$ and $\varphi_k(L \cap V(B_2)) = \{1, \dots, \left\lfloor \frac{k}{2} \right\rfloor\}$.
2. For every leaf λ of $T(k)$, $\varphi_k(V(P(\lambda)) \setminus \{r\}) = [k]$.
3. Each label $i \in [k]$ occurs at most $\lceil \log k \rceil + 1$ times in the labeling of $T(k)$.




Figure 3.2

Proof. We proceed by induction on k . Property 1 follows directly from the recursive definition of the labeling φ_k .

For 2, let λ be a leaf and let the (ordered) vertices of $P(\lambda)$ be $r, p_1, \dots, p_k = \lambda$. Suppose that $\lambda \in V(B_i)$. Then $P(\lambda) := P_i \cup P'$, where P_i is a main path of $T(k)$ and P' is the path in B_i going from the root of B_i to λ . Property 2 now follows by induction and the definition of φ_k .

For 3, first suppose that the label i is in $[k_1]$. Then i appears exactly once in the labeling of the main path P_1 of $T(k)$, it does not figure in the labeling of the nodes $V(P_2) \cup (V(B_1) \setminus V(P_1))$, and, by the inductive step, it occurs $\lceil \log \lceil \frac{k}{2} \rceil \rceil + 1 = \lceil \log k \rceil$ times in $\varphi_k(B_2)$. The thesis follows. A similar argument settles the remaining case $i \in [k] \setminus [k_1]$. \square

Henceforth, we simplify notation and denote the labeling φ_k of $T(k)$ as φ . We now recall some notation from the previous section. Let q be a prime power and S_q be the edge vs stable set incidence matrix of $\mathcal{S}(q)$.

A maximal rectangle $R = (\mathcal{P}_R, \mathcal{L}_R)$ is *centered* if $|\mathcal{L}_R| \geq 2$ and there is a point $c \in \mathcal{P} \setminus \mathcal{P}_R$ such that c is incident to all lines in \mathcal{L}_R . We call c the *center* of R . Note that the center is unique and its existence implies that $|\mathcal{L}_R| \leq q + 1$.

One way to create centered rectangles is as follows. Let ℓ be a line, c be a point on ℓ , and $Y \subseteq N(\ell)$ with $c \in Y$. We let $\boxed{c, \ell, Y}$ be the centered rectangle $R = (\mathcal{P}_R, \mathcal{L}_R)$ where $\mathcal{P}_R = N(\ell) \setminus Y$ and $\mathcal{L}_R = N(c)$. Note that a special entry of the form $(p\ell, S(X))$ is covered by the centered rectangle $\boxed{c, \ell, Y}$ if and only if $p \notin Y$ and $c \in X \subseteq Y$.

We now fix a line $\ell \in \text{PG}(2, q)$ and let $N(\ell) = \{p_1, \dots, p_{q+1}\}$. We will use the labeling φ of $T(q+1)$ to provide a collection of centered rectangles that cover all special entries of the form $(p\ell, S(X))$. Recall that for a vertex v of $T(q+1)$, $P(v)$ denotes the path in $T(q+1)$ from r to v . If v is

Chapter 3. On the extension complexity of the stable set polytope of bipartite graphs

neither the root nor a leaf of $T(q+1)$, we define

$$Y(v) := \{p_{\varphi(u)} \mid u \text{ is a non-root vertex of } P(v)\}.$$

Lemma 3.21. Fix a line $\ell \in \text{PG}(2, q)$ and let $N(\ell) = \{p_1, \dots, p_{q+1}\}$. Let \mathcal{R}_ℓ be the collection of all centered rectangles $\boxed{p_{\varphi(v)}, \ell, Y(v)}$ where v ranges over all non-root, non-leaf vertices of $T(q+1)$. Then every special entry (e, S) with ℓ incident to e is covered by some rectangle $R \in \mathcal{R}_\ell$.

Proof. Let $(p_i \ell, S(X))$ be such a special entry and let λ be the (unique) leaf of $T(q+1)$ such that $\varphi(\lambda) = i$. Name the vertices of $P(\lambda)$ as $r, u_1, \dots, u_{q+1} = \lambda$ (ordered away from the root).

Define $j = \max\{i \mid p_{\varphi(u_i)} \in X\}$. Since $p_{\varphi(\lambda)} \notin X$, note $j < q+1$. By Lemma 3.20, $X \subseteq Y(u_j)$. Also, by construction, $p_{\varphi(u_j)} \in X$ and $p \notin Y(u_j)$. We conclude that the centered rectangle $\boxed{p_{\varphi(u_j)}, \ell, Y(u_j)}$ covers the special entry $(p_i \ell, S(X))$, as required. \square

By Lemma 3.21, for each line ℓ , there is a set \mathcal{R}_ℓ of $O(q \log q)$ centered rectangles that cover all special entries of the form $(p\ell, S(X))$. By taking the union of all \mathcal{R}_ℓ , we get a cover \mathcal{R} of size $O(nq \log q)$ for all the special entries. To prove the main theorem of this section, we now reduce the size of \mathcal{R} by a factor of q .

Theorem 3.22. There is a set of $O(n \log n)$ centered rectangles that cover all the special entries.

Proof. If $R_1 := \boxed{c, \ell_1, Y_1}, \dots, R_k := \boxed{c, \ell_k, Y_k}$ are centered rectangles with the same center c , we let $\sum_{i=1}^k R_i = R$ be the maximal rectangle with $\mathcal{P}_R = \bigcup_{i=1}^k N(\ell_i) \setminus \bigcup_{i=1}^k Y_i$ and $\mathcal{L}_R = N(c)$. Note that $\sum_{i=1}^k R_i$ is also a centered rectangle with center c .

Claim 3.23. If $R_1 := \boxed{c, \ell_1, Y_1}, \dots, R_k := \boxed{c, \ell_k, Y_k}$ are centered rectangles such that ℓ_1, \dots, ℓ_k are all distinct, then $\sum_{i=1}^k R_i$ covers all special entries covered by $\bigcup_{i=1}^k R_i$.

Subproof. Let $(p\ell, S(X))$ be a special entry covered by some $\boxed{c, \ell_j, Y_j}$. Clearly $c \in X \subseteq Y_j \subseteq \bigcup_{i=1}^k Y_i$. By contradiction, suppose $p \in \bigcup_{i=1}^k Y_i$. Since $p \notin Y_j$, $p \in Y_{j'} \subseteq N(\ell_{j'})$ for some $j' \neq j$. But then $c\ell_j p\ell_{j'}$ is a 4-cycle in $\mathcal{J}(q)$, which is a contradiction. Hence the entry $(p\ell, S(X))$ is also covered by $\sum_{i=1}^k R_i$. \blacksquare

We iteratively use Claim 3.23 to reduce the number of rectangles in our covering \mathcal{R} . For each point c , name the $q+1$ lines through c as $\ell, \ell_1, \dots, \ell_q$, so that among $\mathcal{R}_\ell, \mathcal{R}_{\ell_1}, \dots, \mathcal{R}_{\ell_q}$, the collection \mathcal{R}_ℓ has the most rectangles with center c . Note that, by Lemma 3.20, \mathcal{R}_ℓ contains $O(\log q)$ rectangles with center c .

Fix $i \in [q]$ and for each rectangle $R \in \mathcal{R}_{\ell_i}$ with center c choose a rectangle $f_i(R)$ with center c in \mathcal{R}_ℓ such that $f_i(R) \neq f_i(R')$ if $R \neq R'$. For each $R \in \mathcal{R}_{\ell_i}$ we let

$$f^{-1}(R) = \{R\} \cup \bigcup_{i=1}^q \{R' \in \mathcal{R}_{\ell_i} \mid f_i(R') = R\}.$$

We then remove all rectangles with center c that appear in $\mathcal{R}_\ell, \mathcal{R}_{\ell_1}, \dots, \mathcal{R}_{\ell_q}$ and replace them with all rectangles of the form $\sum_{R' \in f^{-1}(R)} R'$, where R ranges over all rectangles in \mathcal{R}_ℓ with center c . In doing so, we obtain at most $O(\log q) = O(\log n)$ rectangles with center c . Repeating for every $c \in \mathcal{P}$ gives us $O(n \log n)$ rectangles in total. \square

3.6 A connection with rectangle covers of the spanning tree polytope

In the previous section we dealt with the problem of covering special entries of the slack matrix S_q with centered rectangles. While this problem may seem at first very specific, in this section we will show that it is equivalent to the class cover problem, that has been introduced by Kaibel and Averkov in [60] in the context of finding rectangle covers for the spanning tree polytope. We first define class covers and show its connections with Section 3.5, as well as a lower bound on the size of class covers which uses the same technique as in the proof of Theorem 3.11. Then, we describe how this problem is related to the extension complexity of the spanning tree polytope and we conclude with some open questions on the topic.

3.6.1 The class covering problem

Let $n \geq 2$ be a natural number. Given $t \in [n]$ and $\emptyset \neq S \subset [n] - t$, we call a couple (t, S) a *class*.

Definition 3.24. A *class cover* of $[n]$ is a family \mathcal{C} of classes, such that for every $s \in [n]$ and $\emptyset \neq X \subset [n] - s$ there is a class $(t, S) \in \mathcal{C}$ with $s \in S$, $t \in X$ and $S \cap X = \emptyset$ (in this case we say that the class *covers* (s, X)). We define $\kappa(n)$ to be the minimum size of a class cover of $[n]$.

A natural, informal interpretation of the problem is as follows: a class (t, S) is formed by a teacher t and a set S of students, and a subject is identified with the set X of people who can be a teacher for the subject. Then a class cover is a family of classes such that every student s can “learn” any subject X through an appropriate class, where we require that none of the students knows the subject ($S \cap X = \emptyset$).

As an example, consider the family $\mathcal{C} = \{(t, \{s\}) : s, t \in [n], s \neq t\}$. It is easy to check that \mathcal{C} is a class cover, which implies $\kappa(n) < O(n^2)$.

The following has been shown by Kaibel and Averkov [60]:

Theorem 3.25. $\kappa(n) = O(n \log n)$.

The reader might notice that this bound is the same as in Theorem 3.22. This is not a coincidence, in fact in Section 3.6.2 we will derive Theorem 3.25 as a corollary of the results in Section 3.5.

Chapter 3. On the extension complexity of the stable set polytope of bipartite graphs

We now show that the bound is asymptotically tight, i.e. $\kappa(n) = \Theta(n \log n)$. The proof uses a very similar technique as in Theorem 3.11. This is joint work with Jana Cslovjcek.

Theorem 3.26. *Let \mathcal{C} be a class cover for $[n]$, then \mathcal{C} has size $\Omega(n \log n)$.*

Proof. To every couple (s, X) with $\emptyset \neq X \subset [n] - s$, we assign the following weight:

$$w(s, X) = \frac{1}{|X| \binom{n-1}{|X|}}.$$

Recall that a class (t, S) covers (s, X) if $s \in S$, $t \in X$ and $S \cap X = \emptyset$. We refer to the weight ‘covered’ by the class as the sum of the weight of all couples (s, X) covered by the class. We will show that the total weight of all the couples (s, X) is at least $cn \log n$ for some positive constant c , which means that the classes in \mathcal{C} have to cover a weight of at least $cn \log n$. On the other hand we show that a single class covers a weight of 1, which implies that we need at least $cn \log n$ different classes to cover all couples (s, X) .

First we calculate the total weight of all couples (s, X) :

$$\sum_{(s,X)} w(s, X) = \sum_{s \in [n]} \sum_{X \subset [n] \setminus s} \frac{1}{|X| \binom{n-1}{|X|}} = \sum_{s \in [n]} \sum_{k=1}^{n-1} \binom{n-1}{k} \frac{1}{k \binom{n-1}{k}} = n \cdot \sum_{k=1}^{n-1} \frac{1}{k} \geq cn \log n,$$

where again $c > 0$ is an appropriate constant.

We now calculate the weight covered by a single class (t, S) :

$$\begin{aligned} \sum_{\substack{(s,X): S \cap X = \emptyset, \\ s \in S, t \in X}} w(s, X) &= \sum_{s \in S} \sum_{t \in X \subset [n] \setminus S} \frac{1}{|X| \binom{n-1}{|X|}} = \sum_{s \in S} \sum_{k=1}^{n-|S|} \binom{n-|S|-1}{k-1} \frac{1}{k \binom{n-1}{k}} \\ &= |S| \sum_{k=1}^{n-|S|} \frac{(n-1-|S|)!}{(k-1)!(n-k-|S|)!} \frac{k!(n-1-k)!}{k(n-1)!} \\ &= \sum_{k=1}^{n-|S|} \frac{(n-1-|S|)!|S|!}{(n-1)!} \frac{(n-1-k)!}{(|S|-1)!(n-k-|S|)!} \\ &= \sum_{k=1}^{n-|S|} \frac{1}{\binom{n-1}{|S|}} \binom{n-1-k}{|S|-1} = \frac{1}{\binom{n-1}{|S|}} \sum_{k=1}^{n-|S|} \binom{n-1-k}{|S|-1} = \frac{1}{\binom{n-1}{|S|}} \binom{n-1}{|S|} = 1 \end{aligned}$$

where the second to last equation comes from Lemma 3.6.

□

3.6.2 Special entries, centered rectangles and class covers

We recall that in Section 3.5 we defined centered rectangles and described a cover of all special entries of S_q (slack matrix of $\mathcal{J}(q)$) with centered rectangles. To obtain this we first, for every

line $\ell \in \text{PG}(2, q)$, gave a covering \mathcal{R}_ℓ of all the special entries (e, S) with ℓ incident to e , and then merged the covers in one. We now show how this problem is related to the class covering problem.

Theorem 3.27. *Let q be a prime power and S_q be the edge vs stable set incidence matrix of $\mathcal{J}(q)$. For a fixed line $\ell \in \text{PG}(2, q)$, let CR_ℓ be the set of centered rectangles of S_q of the form $R = (\mathcal{P}_R, \mathcal{L}_R)$ with $\mathcal{P}_R = N(\ell) \setminus Y$ and $\mathcal{L}_R = N(c)$ for some $c \in Y \subset N(\ell)$, and SE_ℓ be the set of special entries (e, S) of S_q with ℓ incident to e . Then the problem of covering entries in SE_ℓ with rectangles in CR_ℓ is equivalent to finding a class cover for $[n]$, with $n = q + 1$.*

Proof. Each rectangle in CR_ℓ can be described by the couple (c, Y) , with $c \in Y \subset N(\ell)$, and each special entry in SE_ℓ is given by a couple (p, X) where $p \in N(\ell)$ and $X \subseteq N(\ell) \setminus \{p\}$. Now, recall that a centered rectangle (c, Y) covers a special entry (p, X) if and only if

- $p \in N(\ell) \setminus Y$, and
- $c \in X \subseteq Y$.

Mapping the points in $N(\ell)$ to $[n]$ (recall that $\mathcal{J}(q)$ is $q + 1$ -regular), we can consider (p, X) as a pair of a student and a subject and $(c, [n] \setminus Y)$ as a class, and the equivalence becomes now clear: (p, X) is covered by the class $(c, [n] \setminus Y)$ if and only if $p \in [n] \setminus Y$, $c \in X$, $X \cap [n] \setminus Y = \emptyset$, i.e. $X \subset Y$. \square

As a consequence of Theorem 3.27, we can use the construction from Section 3.5 to obtain a class cover for $[n]$ of size $\Theta(n \log n)$, proving Theorem 3.25. In particular, given the tree $T(n)$ with labeling $\varphi: V(T(n)) \rightarrow [n]$ defined in Section 3.5, let $S(v) = [n] \setminus \varphi(V(P(v)))$ for any non-root, non-leaf vertex v of $T(n)$, i.e. $S(v)$ is the set of labels on any path from v to a leaf of $T(n)$, such that the path does not go through the root of $T(n)$. Notice that, thanks to Lemma 3.20, part 2, $S(v)$ does not depend on which path we choose. Then it is easy to check that the set $\mathcal{C} = \{(\varphi(v), S(v)) : v \in V(T(n))\}$ is a class cover for $[n]$ of size $|V(T(n))| = \Theta(n \log n)$. For an example, consider the tree $T(3)$ and its labeling, as described in Figure 3.2a. The resulting class cover is formed by the following classes: $(1, \{2\})$, $(1, \{2, 3\})$, $(2, \{1\})$, $(2, \{3\})$, $(3, \{1, 2\})$. Class covers constructed this way are asymptotically optimal thanks to Theorem 3.26.

3.6.3 Spanning tree polytopes

As already mentioned at the beginning of this chapter, obtaining more precise bounds on the extension complexity of the spanning tree polytope $\text{STP}(G)$ is an open problem, on which almost no progress has been done since the extended formulation of Wong [99] and Martin [79]. For $G = K_n$, we have that $\Omega(n^2) = \text{xc}(\text{STP}(K_n)) = O(n^3)$. In [65], Khoshkhah and Theis prove that using the rectangle covering bound defined in Section 3.2 cannot help to improve the lower bound on $\text{xc}(\text{STP}(K_n))$ apart from a logarithmic factor.

Theorem 3.28. [65] Let S_n^{STP} be the slack matrix of $\text{STP}(K_n)$, then $\text{rc}(S_n^{\text{STP}}) \leq O(n^2 \log n)$.

In the original proof, the bound is obtained through a connection with communication complexity and the rectangle cover is only implicitly given. In [60] Kaibel and Averkov give an explicit rectangle cover proving Theorem 3.28, using class covers.

For the complete graph $K_n = (V, E)$, we have:

$$\begin{aligned} \text{STP}(K_n) = \{x \in \mathbb{R}^E : & \sum_{e \in E(U)} x_e \leq |U| - 1 \quad \forall U \subset V, |U| > 1 \\ & x_e \geq 0 \quad \forall e \in E \\ & \sum_{e \in E} x_e = n - 1 \quad \}. \end{aligned} \quad (3.1)$$

We only consider the submatrix $S_n^{\text{STP}'}$ of S_n^{STP} corresponding to the first set of inequalities, as the rest can be trivially covered by $O(n^2)$ rectangles (we refer to [65] for further details). For $U \subset V$ and a tree T , $S_n^{\text{STP}'}$ has as entry (U, T) the number of connected components in $(U, T(U))$ minus 1. In particular, an entry (U, T) is non-zero if $T(U)$ is not connected.

Theorem 3.29. [60] For $n \geq 2$, $\text{rc}(S_n^{\text{STP}'}) \leq n \cdot \kappa(n)$

Proof. For A, B disjoint non-empty subsets of $V = [n]$ and T a spanning tree of K_n , we say that (A, B) is a T -disconnecter if there exist $u, v \in A$ such that the path from u to v on T has a node in B . The following can be easily seen to be a rectangle of $S_n^{\text{STP}'}$:

$$R(A, B) = \{U : A \subseteq U \subseteq [n] \setminus B\} \times \{T \text{ spanning tree} : (A, B) \text{ is a } T\text{-disconnecter}\}$$

Now, let \mathcal{C} be a class cover for $[n]$. We will show that the set

$$\mathcal{R} = \{R(\{v, t\}, S) : (t, S) \in \mathcal{C}, v \in [n] \setminus (S \cup \{t\})\}$$

is a rectangle covering for $S_n^{\text{STP}'}$, which concludes the proof.

We need to show that for every non zero entry (U, T) of the slack matrix there exists $R(A, B) \in \mathcal{R}$ such that $(U, T) \in R(A, B)$. Since the entry (U, T) is non zero we have that $T(U)$ is disconnected. Let u, v be in two different connected components of U . Then there is a vertex w on the path between u and v which is not in U . Since \mathcal{C} is a class cover, in correspondence of the student-subject pair (w, U) there is a class $(t, S) \in \mathcal{C}$ such that $t \in U \subseteq [n] \setminus S$, $w \in S$. Now, we have that w is on the path of T between t and v , or between t and u : assume without loss of generality that the first case holds. Then $(U, T) \in R(\{t, v\}, S)$: indeed, $\{t, v\} \subseteq U \subseteq [n] \setminus S$ and $(\{t, v\}, S)$ is a T -disconnecter. \square

Now Theorem 3.28 follows directly from Theorem 3.25.

3.6. A connection with rectangle covers of the spanning tree polytope

As we have shown, the class covering problem has been successfully used for constructing rectangle covers in two apparently unrelated settings. Whether it could be applied to investigate the extension complexity of other problems is a fascinating question. Moreover, Theorem 3.28 does not rule out that a $\Omega(n^2 \log n)$ lower bound for $\text{xc}(\text{STP}(K_n))$ could be proved by rectangle covering techniques. While such a bound is not directly implied by Theorem 3.11 and 3.26, it might be possible to attack the problem with the method used to prove those theorems.

4 Slack matrices of 2-level polytopes: recognition and decomposition

4.1 Introduction

Slack matrices are interesting mathematical objects that carry information on the vertex-facet adjacency structure of a polytope and on its extension complexity, as we already discussed in Chapter 1. In an attempt to shed light on the properties of such matrices, in [43] Gouveia et al. give some geometric characterizations of slack matrices and study the following problem: given a non-negative matrix M , can we decide in polynomial time whether M is the slack matrix of some polytope? The authors prove that this problem, called *slack matrix recognition problem*, is equivalent to the following: given $P = \{x \in \mathbb{R}^d : Ax \leq b\}$, and $Q = \text{conv}\{v_1, \dots, v_n\} \subseteq P$, decide whether $Q = P$. This is known as the Polyhedral Verification problem, a central problem in computational geometry whose complexity is unknown ([61]). It is therefore natural to try to find polynomial algorithms for recognizing restricted classes of slack matrices. In particular we focus on 0/1 slack matrices: those are exactly the slack matrices of 2-level polytopes. It is not clear how to approach such a problem since, as argued in the introduction, we are far from a complete understanding of 2-level polytopes and their slack matrices. On the other hand, progress on this problem would most likely advance our knowledge on 2-level polytopes. In this chapter we will describe some algorithmic results on the recognition of certain sub-classes of 0/1 slack matrices. We will also define some operations on (slack matrices of) polytopes that preserve 2-levelness, which can be seen as modifications of the operation of cartesian product (see Section 1.1). Our study of such operations is motivated by their application in recognizing slack matrices of 2-level matroid polytopes. We recall (see Theorem 2.21) that matroids whose base polytope is 2-level arise from uniform matroids by applying 1-sums and 2-sums (we refer to Section 2.4 for the relevant definitions). By studying such operations in the context of slack matrices, we are able to recognize slack matrices of 2-level matroid polytopes in polynomial time. Moreover, we define the more general operation of k -sum of slack matrices and we investigate its properties. This might lead to decomposition results that would have high impact on the open question on 2-level polytopes.

Contribution and organization.

The chapter is organized as follows:

- after introducing some basic properties of slack matrices, we define the operations of 1-sum and k -sum, for $k \geq 2$, in Section 4.2, prove some of their properties and present some experimental evidence for the relevance of such operations in the context of 2-level polytopes.
- In Section 4.3 we investigate slack matrices of stable set polytopes of perfect graphs: we characterize when such matrices are k -sums and also provide a simple algorithm for recognizing them.
- We then study the problem of recognizing matrices that are k -sums, and provide efficient algorithms in Section 4.4. At the core of such algorithms is a connection between the 1-sum operation and mutual entropy, a function used in information theory.
- The results obtained are applied to the recognition of slack matrices of 2-level matroid polytopes in Section 4.4.4.
- To conclude, in Section 4.5 we describe an alternative approach for the latter problem, which might be applicable to a larger class of matroid polytopes.

4.1.1 Preliminaries

For the definition of slack matrix, we refer the reader to Definition 1.3. In [43], the authors characterize slack matrices of polytopes. For a matrix S , we denote the collection of column vectors of S by $\text{col}(S)$.

Theorem 4.1 ([43]). *For a nonnegative matrix S in $\mathbb{R}^{m \times n}$ of rank at least 2, the following statements are equivalent:*

1. S is the slack matrix of a polytope;
2. $\text{conv}(\text{col}(S)) = \text{aff}(\text{col}(S)) \cap \mathbb{R}_+^m$;
3. $\mathbb{R}^m \cdot S \cap \mathbb{R}_+^n = \mathbb{R}_+^m \cdot S$ and $\mathbf{1} \in \mathbb{R}^m \cdot S$ holds;
4. $S \cdot \mathbb{R}^n \cap \mathbb{R}_+^m = S \cdot \mathbb{R}_+^n$ and $\mathbf{1} \in \mathbb{R}^m \cdot S$ holds.

Throughout the chapter, we will assume that all the matrices we deal with are of rank at least 2, so to apply Theorem 4.1 directly. We also recall the following useful fact:

Lemma 4.2. [43] *Let S be a slack matrix of a polytope P , then P is affinely isomorphic to $\text{conv}(\text{col}(S))$. In addition, we have $\dim(P) = \text{rk}(S) - 1$.*

Recall that the slack matrix of a polytope P is not unique, as it depends on the given horizontal and vertical representations of P . We say that a slack matrix is *non-redundant* if the two representations are, i.e. if the matrix has exactly as many rows as P has facets and as many

columns as P has vertices. Non-redundant slack matrices do not contain two identical rows or columns, nor rows or columns which are all zeros, or all non-zeros. In light of Lemma 4.2, using linear programming one can efficiently check whether some columns are redundant (i.e. are contained in the convex hull of the others). To check whether a (not all-zero) row is redundant it suffices to check whether its set of zeros (i.e. the vertices lying on the corresponding face) is maximal.

As mentioned in Section 1, a polytope is 2-level if and only if it has a 0/1 slack matrix. For 0/1 matrices, checking for non-redundancy does not require linear programming: indeed, the columns are always non-redundant as they can be seen as vertices of a hypercube, and for the rows it again suffices to look at the set of zeros of each row. Hence, for simplicity, when dealing with the slack matrix recognition problem we can always assume that our candidate slack matrix is non-redundant. However, in the following we will sometimes assume that S contains a certain row corresponding to a face of P , even if such row may be redundant. In particular we observe that if r is a non-redundant row of S , the row $\tilde{1} - r$ represents a face of P hence it can be added to S without changing the fact that it is a slack matrix of P . This is a consequence of the 2-level property: indeed, the vertices that do not lie on facet corresponding to r are contained in a single hyperplane, hence they form a face of P (not necessarily a facet), whose slack is given exactly by $\tilde{1} - r$.

4.2 k -sums of slack-matrices

Given two non-empty matrices S_1 and S_2 , we define the operations of 1- and 2-sum of S_1 and S_2 , which generalizes to every $k \geq 1$, and we show that these operations essentially preserve the property of being a slack matrix for 0/1 matrices. For this reason we will only deal with 0/1 matrices, even though some of the definitions and results hold for more general settings, for instance for matrices with real entries. The definition of k -sum is similar to the notion of glued product appearing in [67, 77], but it has been defined independently from it. [67] contains results analogous to Lemma 4.11 and Corollary 4.12, with the difference that the operation of glued product is defined on polytopes, while our k -sum is defined on general matrices.

4.2.1 1-sums

Definition 4.3 (1-sum). The 1-sum of $S_1 \in \{0, 1\}^{m_1 \times n_1}$ and $S_2 \in \{0, 1\}^{m_2 \times n_2}$ is the matrix S whose set of columns is obtained concatenating every column of S_1 with every column of S_2 . More precisely, for $i = 1, \dots, n_1 \cdot n_2$, $S^i = \begin{bmatrix} S_1^j \\ S_2^k \end{bmatrix}$, with $i = (j-1)n_2 + k$, and where M^ℓ denotes the ℓ -th column of matrix M . We write $S = S_1 \oplus_1 S_2$ or simply $S = S_1 \oplus S_2$. Notice that S has $m_1 + m_2$ rows and $n_1 \cdot n_2$ columns.

Example 4.4. Here follow some examples of 1-sums:

$$\begin{bmatrix} 1 & 0 \end{bmatrix} \oplus \begin{bmatrix} 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{bmatrix} \oplus \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 1 \end{bmatrix}.$$

For the purpose of recognizing slack matrices, permutations of columns and rows are not relevant: the property of being a slack matrix is preserved when two rows or two columns are swapped. Hence we call two matrices *isomorphic* if one can be obtained from the other by permuting rows and columns. We say that a matrix S is a 1-sum if there are two matrices S_1, S_2 (each with strictly less rows than S) such that S is isomorphic to $S_1 \oplus S_2$ (we will often abuse notation and simply write $S = S_1 \oplus S_2$). We call a matrix *irreducible* if it is not a 1-sum.

- Remark 4.5.**
1. Notice that S has a pair of identical columns if and only if the same is true for either S_1 or S_2 or both. Similarly, S has a pair of identical, non-constant rows (i.e. containing at least a 0 and a 1) if and only if at least one of S_1, S_2 does.
 2. Let R, R_1, R_2 be the set of rows of S, S_1, S_2 respectively. Then there is a natural bijection between R and $R_1 \cup R_2$ (where the union is disjoint): up to element permutation, each row in R is obtained by concatenating a row in R_1 with itself n_2 times or a row in R_2 with itself n_1 times. Hence, R_1, R_2 induce a partition of R , and we say that S is a 1-sum with respect to R_1, R_2 . In particular, if S is a 1-sum with respect to R_1, R_2 , then $S|_{R_1}$ (restricted to the rows R_1) is made of n_2 copies of S_1 (possibly permuted), and similarly $S|_{R_2}$ is made of n_1 copies of S_2 . Notice that the choice of such S_1, S_2 is not necessarily unique, unless S has no repeated columns: in this case S_1 is obtained by keeping each column of $S|_{R_1}$ exactly once, and similarly for S_2 .
 3. The above remarks can be extended to the 1-sum of three or more matrices.

We will prove now that the 1-sum operation preserves the property of being a slack matrix. In particular if S_i is the slack matrix of some polytope P_i , $i = 1, 2$, then S is the slack matrix of the Cartesian product $P_1 \times P_2 = \{(x, y) \in \mathbb{R}^{d_1+d_2} : x \in P_1, y \in P_2\}$, where d_i denotes the dimension of polytope P_i for $i = 1, 2$.

Lemma 4.6. *Let $S \in \{0, 1\}^{m \times n}$ and let $S_i \in \{0, 1\}^{m_i \times n_i}$ for $i = 1, 2$ such that $S = S_1 \oplus S_2$. Then S is a slack matrix of a polytope P if and only if there exist polytopes P_1, P_2 such that S_i is a slack matrix of P_i for $i = 1, 2$ and $P \cong P_1 \times P_2$.*

Proof. From Theorem 4.1 and Lemma 4.2, it follows that the thesis is equivalent to proving that

$$\text{aff}(\text{col}(S_1 \oplus S_2)) = \text{aff}(\text{col}(S_1)) \times \text{aff}(\text{col}(S_2))$$

and

$$\text{conv}(\text{col}(S_1 \oplus S_2)) = \text{conv}(\text{col}(S_1)) \times \text{conv}(\text{col}(S_2)).$$

The same proof works for both statements, hence we only prove the first one. The inclusion from left to right is easy. We prove ‘ \supseteq ’. Take $p = (x, y)$ a point in $\text{conv}(\text{col}(S_1)) \times \text{conv}(\text{col}(S_2))$. Then x is an affine combination $\sum_i \lambda_i x_i = x$ of the columns x_i of S_1 , $i = 1, \dots, n_1$ and, similarly, y is an affine combination $\sum_j \mu_j y_j = y$ of the columns y_j of S_2 , $j = 1, \dots, n_2$. We deduce that $p = (x, y) = \sum_i \sum_j \lambda_i \mu_j (x_i, y_j)$, and $\sum_{i,j} \lambda_i \mu_j = \sum_i \lambda_i \sum_j \mu_j = \sum_i \lambda_i = 1$. Moreover, if $\mu_j \geq 0, \lambda_i \geq 0$ for any i, j , then the multipliers are all non-negative, proving the second statement. \square

The following is an immediate consequence of Lemma 4.6 and Lemma 4.2.

Corollary 4.7. *Let S, S_1, S_2 be slack matrices satisfying $S = S_1 \oplus S_2$. Then $\text{rk}(S) = \text{rk}(S_1) + \text{rk}(S_2) - 1$.*

4.2.2 2-sums and k -sums

In this section we define a more general operation of k -sum, for $k \geq 2$. We first treat the case $k = 2$. As before, let $S_1 \in \{0, 1\}^{m_1 \times n_1}$ and $S_2 \in \{0, 1\}^{m_2 \times n_2}$, and let x_1, y_1 be rows of S_1, S_2 , respectively. We call x_1, y_1 *special* rows. For any matrix M and row r of M , we denote by $M - r$ the matrix obtained from M by removing row r . The row x_1 determines a partition of the columns of $S_1 - x_1$ according to its 0 and 1 entries: we construct the submatrices S_1^0, S_1^1 as

$$S_1^a := ((S_1 - x_1)_{ij} : x_{1j} = a) \quad \text{for } a = 0, 1. \quad (4.1)$$

$$S_1 = \left(\begin{array}{c|c} S_1^0 & S_1^1 \\ \hline 0 \cdots 0 & 1 \cdots 1 \end{array} \right) \leftarrow x_1$$

Figure 4.1

Similarly, y_1 induces a partition of $S_2 - y_1$ into S_2^0, S_2^1 . Here we assume that none of $S_1^0, S_1^1, S_2^0, S_2^1$ is empty, i.e. that each of x_1, y_1 contains at least a 0 and a 1.

We consider the matrix S_1 with special row x_1 as a pair (S_1, x_1) , and matrix S_2 with special row y_1 as (S_2, y_1) .

Definition 4.8 (2-sum). With the previous notations, the 2-sum S of (S_1, x_1) and (S_2, y_1) ,

denoted by $S = (S_1, x_1) \oplus_2 (S_2, y_1)$, is:

$$S = (S_1, x_1) \oplus_2 (S_2, y_1) := \left(\begin{array}{c|c} S_1^0 \oplus S_2^0 & S_1^1 \oplus S_2^1 \\ \hline 0 \cdots 0 & 1 \cdots 1 \end{array} \right)$$

Similarly as before, we say that S is a 2-sum if there exist matrices S_1, S_2 (each with less rows and columns than S) and rows x_1 of S_1 , y_1 of S_2 , such that S is isomorphic to $(S_1, x_1) \oplus_2 (S_2, y_1)$. Again, we will abuse notation and write $S = (S_1, x_1) \oplus_2 (S_2, y_1)$.

We will now extend Lemma 4.6 to the 2-sum operation. However, the 2-sum does not behave as well as the 1-sum in terms of slack matrix, as the following example shows. Assume S_1, S_2 are non-redundant slack matrices such that S_1 has two opposite rows $x_1, x_2 = \mathbf{1} - x_1$, and similarly S_2 has two opposite rows $y_1, y_2 = \mathbf{1} - y_1$. Let $S = (S_1, x_1) \oplus_2 (S_2, y_1)$, in the next lemma we will prove that S is a slack matrix as well. It is not hard to see that S has two identical rows (in correspondence of x_2, y_2), hence we can delete one of them and obtain S' , which is obviously still a slack matrix, and satisfies $S' = (S_1, x_1) \oplus_2 (S'_2, y_1)$ where $S'_2 = S_2 - y_2$. But now S'_2 does not need to be a slack matrix anymore. This could in principle cause problems since when decomposing a slack matrix into a 2-sum we might obtain with factors that are not slack matrices. However we will see that this example is the only exception, one for which there is a simple fix: one just adds to each factor the row which is opposite to the special row. As observed in Section 4.1.1, such opposite rows can be safely added while preserving the property of being a slack matrix (and in practice, they need to be added only if they are not dominated by some other row).

The proof of the following lemma is omitted as it is a special case of Lemma 4.11.

Lemma 4.9. *Let $S \in \{0, 1\}^{m \times n}$ and let $S_i \in \{0, 1\}^{m_i \times n_i}$ for $i = 1, 2$ such that $S = (S_1, x_1) \oplus_2 (S_2, y_1)$ for some row x_1 of S_1 , y_1 of S_2 . Then the following hold:*

1. *If S_1, S_2 are slack matrices, then S is a slack matrix.*
2. *If S is a slack matrix, let $S'_1 = S_1 + (\mathbf{1} - x_1)$ (i.e. S_1 with one additional row that is opposite to x_1), and similarly let $S'_2 = S_2 + (\mathbf{1} - y_2)$. Then S'_1, S'_2 are slack matrices.*

We now define the general operation of k -sum, $k \geq 2$. Similarly as before, we consider two 0/1 matrices S_1, S_2 each with $k - 1$ special rows, x_1, \dots, x_{k-1} , y_1, \dots, y_{k-1} respectively, such that no column of S_1 has more than one entry corresponding to an x_i equal to 1, and similarly for S_2 . For $(a_1, \dots, a_{k-1}) \in \{(0, \dots, 0), (1, \dots, 0), \dots, (0, \dots, 1)\}$, $S_1 - \{x_1, \dots, x_{k-1}\}$ is partitioned into submatrices $S_1^{a_1 \dots a_{k-1}}$ containing all the columns of S_1 whose entries in correspondence of the special rows are a_1, \dots, a_{k-1} , and similarly for $S_2^{a_1 \dots a_{k-1}}$. We assume again that all submatrices $S_1^{a_1 \dots a_{k-1}}, S_2^{a_1 \dots a_{k-1}}$ are non-empty.

Definition 4.10 (k -sum). With the previous notations, the k -sum S of the matrices S_1 with special rows (x_1, \dots, x_{k-1}) and S_2 with special rows (y_1, \dots, y_{k-1}) , denoted by $(S_1, (x_1, \dots, x_{k-1})) \oplus_k$

$(S_2, (y_1, \dots, y_{k-1}))$, is defined to be:

$$S = (S_1, (x_1, \dots, x_{k-1})) \oplus_k (S_2, (y_1, \dots, y_{k-1})) := \left(\begin{array}{c|c|c|c} S^{0\dots 0} & S^{0\dots 1} & \dots & S^{1\dots 0} \\ 0\dots 0 & 0\dots 0 & \dots & 1\dots 1 \\ \vdots & \vdots & \vdots & \vdots \\ 0\dots 0 & 1\dots 1 & \dots & 0\dots 0 \end{array} \right) \quad (4.2)$$

where, for every $(a_1, \dots, a_{k-1}) \in \{(0, \dots, 0), (1, \dots, 0), \dots, (0, \dots, 1)\}$, $S^{a_1\dots a_{k-1}} := S_1^{a_1\dots a_{k-1}} \oplus S_2^{a_1\dots a_{k-1}}$.

Similarly as before, we say that S is a k -sum if there exist matrices S_1, S_2 (each with less rows and columns than S) and rows x_1, \dots, x_k of S_1 , y_1, \dots, y_k of S_2 , such that S is isomorphic to $(S_1, (x_1, \dots, x_{k-1})) \oplus_k (S_2, (y_1, \dots, y_{k-1}))$. Again, we will abuse notation and write $S = (S_1, x_1, \dots, x_{k-1}) \oplus (S_2, y_1, \dots, y_{k-1})$. The following lemma is a generalization of Lemma 4.9.

Lemma 4.11. *Let $S \in \{0, 1\}^{m \times n}$ and let $S_i \in \{0, 1\}^{m_i \times n_i}$ for $i = 1, 2$ such that $S = (S_1, x_1, \dots, x_{k-1}) \oplus_k (S_2, y_1, \dots, y_{k-1})$ for some special rows (x_1, \dots, x_{k-1}) of S_1 , and (y_1, \dots, y_{k-1}) of S_2 .*

1. *If S_1, S_2 are slack matrices, then S is a slack matrix.*
2. *If S is a slack matrix, let $S'_1 = S_1 + (\mathbf{1} - x_1 - \dots - x_{k-1})$, and construct S'_2 similarly. Then S'_1, S'_2 are slack matrices.*

Proof. 1. Let $P_i := \text{conv}(\text{col}(S_i)) \subseteq \mathbb{R}^{m_i}$ for $i = 1, 2$. Without loss of generality, x_1, \dots, x_{k-1} can be assumed to be the first $k-1$ rows of S_1 , and similarly for y_1, \dots, y_{k-1} and S_2 . Hence, for a point $x \in \mathbb{R}^{m_1}$, we overload notation and denote by x_i the i -th coordinate of x , and similarly for $y \in \mathbb{R}^{m_2}$.

By Lemma 4.6, we have that S is a submatrix of a slack matrix of $P_1 \times P_2 \cap H$, where H is the hyperplane defined by the equations $x_1 = y_1, \dots, x_{k-1} = y_{k-1}$. Proving that S is a slack matrix is equivalent to showing that intersecting $P_1 \times P_2$ with H does not create any new vertex, i.e. that $P_1 \times P_2 \cap H \subseteq \text{conv}(\text{col}(S))$.

Consider a point $p = (x^*, y^*) \in P_1 \times P_2 \cap H \subseteq \mathbb{R}^{m_1} \times \mathbb{R}^{m_2}$, for some $x^* \in P_1$ and $y^* \in P_2$. For a density argument, it suffices to show that (x^*, y^*) is in $\text{conv}(\text{col}(S))$ when (x^*, y^*) has all rational coordinates. Then x^* is a convex combination of the vertices of P_1 and y^* is a convex combination of the vertices of P_2 (which are columns of S_1, S_2 respectively), where the coefficients are all rational:

$$x^* = \sum_{i=1}^{n_1} \lambda_i v_i \quad \text{and} \quad y^* = \sum_{j=1}^{n_2} \mu_j w_j \quad \text{with} \quad \sum_i \lambda_i = \sum_j \mu_j = 1, \quad \lambda_i, \mu_j \in \mathbb{Q}_+.$$

Then there exists a positive integer K such that $K\lambda_i \in \mathbb{N}$ and $K\mu_j \in \mathbb{N}$ for every $i \in [n_1]$ and every $j \in [n_2]$. Moreover, $K = K \sum_i \lambda_i = K \sum_j \mu_j$.

Let us partition the set of vertices of P_1 that occur in the convex combination into k subsets, according to the entries corresponding to the special rows (x_1, \dots, x_{k-1}) . For

$\ell = 1, \dots, k-1$, let V_1^ℓ be the set of v_i 's with ℓ -th coordinate equal to 1, and let V_1^0 be the set of the remaining v_i 's (i.e. with the first $k-1$ coordinates all equal to 0, since each v_i has at most one of the first $k-1$ coordinates equal to 1). The sets V_2^0, \dots, V_2^{k-1} are defined similarly. For x^* , we have:

$$x^* = \sum_{v_i \in V_1^0} \lambda_i v_i + \dots + \sum_{v_i \in V_1^{k-1}} \lambda_i v_i.$$

We split in the same way the identity $K = K \sum_i \lambda_i$. Thus $K = \alpha_0 + \dots + \alpha_{k-1}$, where $\alpha_\ell := \sum_{v_i \in V_1^\ell} (K \lambda_i)$ for $\ell = 0, \dots, k-1$. Applying the same reasoning to y^* , we get that $K = \beta_0 + \dots + \beta_{k-1}$, where $\beta_\ell := \sum_{w_j \in V_2^\ell} (K \mu_j)$ for $\ell = 0, \dots, k-1$.

Since the v_i 's and w_j 's are 0/1 vectors, we have that $Kx_\ell^* = \alpha_\ell$ and $Ky_\ell^* = \beta_\ell$ for $\ell = 1, \dots, k-1$. Exploiting the fact that the first $k-1$ coordinates of x^*, y^* are equal, we have that $\alpha_\ell = \beta_\ell$ for $\ell = 1, \dots, k-1$. These identities jointly imply that $\alpha_0 = \beta_0$.

Fix $\ell \in \{0, \dots, k-1\}$. The coefficients α_ℓ, β_ℓ coincide with the number of vectors v_i in V_1^ℓ and w_j in V_2^ℓ when counted with their multiplicity $K\lambda_i$ in the identity of $Kx^* = K \sum_i \lambda_i v_i$ and $K\mu_j$ in $Ky^* = K \sum_j \mu_j w_j$ respectively. Consider then the multiset \bar{V}_1^ℓ containing each vector v_i with multiplicity $K\lambda_i$ and, similarly \bar{V}_2^ℓ containing each vector w_j with multiplicity $K\mu_j$. As $|\bar{V}_1^\ell| = |\bar{V}_2^\ell| = \alpha_\ell$, there exists a bijection Φ_ℓ from the first to the latter. Let $\bar{\Phi}_\ell$ be the truncation of Φ_ℓ excluding coordinates y_1, \dots, y_{k-1} . Hence, the vectors $(v_i, \bar{\Phi}_\ell(v_i))$ are columns of S for every $v_i \in \bar{V}_1^\ell$.

We can now express $p = (x^*, y^*)$ as:

$$(x^*, y^*) = \frac{1}{K} \left(\sum_{v_i \in \bar{V}_1^0} (v_i, \bar{\Phi}_0(v_i)) + \dots + \sum_{v_i \in \bar{V}_1^{k-1}} (v_i, \bar{\Phi}_{k-1}(v_i)) \right)$$

This shows that (x^*, y^*) lies in the convex hull of the columns of S .

2. Let $S = (S_1, x_1, \dots, x_{k-1}) \oplus_k (S_2, y_1, \dots, y_{k-1})$ be a 0/1 slack matrix. We show that $S'_1 = S_1 + (\mathbf{1} - x_1 - \dots - x_{k-1})$ is a slack matrix, the argument for S'_2 is exactly the same. By Theorem 4.1, we have $\text{aff}(\text{col}(S)) \cap \mathbb{R}_+^{m_1} = \text{conv}(\text{col}(S))$, and we will show that the same holds for S'_1 . We use a similar notation than in the first part: we assume that $x_1, \dots, x_{k-1}, \mathbf{1} - x_1 - \dots - x_{k-1}$ are the first k rows of S'_1 , and for $\ell = 1, \dots, k$ we denote by V_1^ℓ the set of columns of S'_1 with ℓ -th coordinate equal to 1 (notice that we do not use V_1^0 anymore for simplicity of notation). Let $x^* \in \text{aff}(\text{col}(S'_1)) \cap \mathbb{R}_+^{m_1}$, one has $x^* = \sum_i \lambda_i v_i = \sum_{v_i \in V_1^1} \lambda_i v_i + \dots + \sum_{v_i \in V_1^k} \lambda_i v_i$, with $\sum_i \lambda_i = 1$. In particular, $x_\ell^* = \sum_{v_i \in V_1^\ell} \lambda_i \geq 0$ for $\ell = 1, \dots, k$ with $x_1^* + \dots + x_k^* = 1$. We now extend x^* to a point $\tilde{x} \in \text{aff}(\text{col}(S))$ as follows: for $\ell = 1, \dots, k$, fix a column u_ℓ of S_2 with the ℓ -th coordinate equal to 1 if $\ell < k$, and with the first $k-1$ coordinates all equal to 0 if $\ell = k$ (such columns exist since by assumption every $S_2^{0 \dots 0}, \dots, S_2^{1 \dots 0}$ is non-empty). Then map each $v_i \in V_1^\ell$ to the column of S consisting of v_i (without its k -th coordinate) followed by u_ℓ (without the coordinates corresponding to y_1, \dots, y_{k-1}) for $\ell = 1, \dots, k$. We denote such column by w_i , for $i = 1, \dots, n_1$, and

let $\tilde{x} = \sum_i \lambda_i w_i$. Now, we claim that $\tilde{x} \in \mathbb{R}_+^m$: indeed, by construction of the w_i 's, every component of \tilde{x} is equal to a component of x^* , or to a sum of a (possibly empty) subset of $\{x_1^*, \dots, x_k^*\}$, according to the corresponding component of the u_ℓ 's. Hence we have that $\tilde{x} \in \text{aff}(\text{col}(S)) \cap \mathbb{R}_+^m = \text{conv}(\text{col}(S))$. We claim that this implies that $x^* \in \text{conv}(\text{col}(S'_1))$: indeed, if $\tilde{x} = \sum_i \mu_i w_i$, with $\mu_i \geq 0$ for $i = 1, \dots, n_1$, and $\sum_i \mu_i = 1$ then it follows that $x^* = \sum_i \mu_i v_i$. This is trivial except for the k -th coordinate: but the latter is equal to $\sum_{v_i \in V_1^k} \mu_i = 1 - \sum_{v_i \in V_1^1} \mu_i + \dots + \sum_{v_i \in V_1^{k-1}} \mu_i = 1 - x_1^* - \dots - x_{k-1}^* = x_k^*$. Hence we conclude that S'_1 is a slack matrix.

□

Corollary 4.12. *Let S, S_1, S_2 be slack matrices of polytopes P, P_1, P_2 respectively, satisfying the hypothesis of Lemma 4.11. Then $xc(P) \leq xc(P_1) + xc(P_2)$.*

Proof. This directly follows from the fact that P is linearly isomorphic to $P_1 \times P_2 \cap H$, where H is the hyperplane defined in the proof of Lemma 4.11. Indeed it is an easy observation that if Q_i is an extension of P_i , i.e. if $P_i = \{x^{(i)} : \exists y^{(i)} : (x^{(i)}, y^{(i)}) \in Q_i\}$ for $i = 1, 2$, then $P_1 \times P_2 = \{(x^{(1)}, x^{(2)}) : \exists (x^{(i)}, y^{(i)}) \in Q_i \text{ for } i = 1, 2\}$.

□

This corollary suggests that decomposition via k -sum can be a useful tool for proving upper bounds on the extension complexity of 2-level polytopes.

We conclude by arguing that, in Lemma 4.11, if S_1 is a slack matrix of P_1 , then the row $r = \mathbf{1} - x_1 - \dots - x_{k-1}$ corresponds to a face of P_1 (and similarly for S_2). Hence, in practice, we only need to add this row if it is not dominated by any of the rows of S_1 . We have that r is a linear combination of the rows of S_1 and is non-negative, hence by Theorem 4.1 it is a conic combination of the rows of S_1 . But this implies that r corresponds to a face of P_1 : for a proof, let $r = \sum_i \lambda_i x_i$, where $\lambda_i \geq 0$ and i ranges over all the row indices of S_1 , and let $a_i z = b_i$ be a hyperplane defining the facet of P_1 corresponding to x_i for every i . Then r corresponds to the face defined by $\sum_i \lambda_i a_i z = \sum_i \lambda_i b_i$. As a last remark, notice that the identity matrix I_k acts as a neutral element for the k -sum of slack matrices: indeed, we have $S_1 \oplus_k I_k = S_1 + r$, and we just argued that S_1 and $S_1 + r$ are essentially the same slack matrix. However, we defined S to be a k -sum if the two factors have strictly less rows and columns than S , thus avoiding this technical issue. In particular it is not hard to see that if both S_1, S_2 are not the identity matrix (which is the slack matrix of a simplex), then their k -sum has rank strictly greater than the rank of S_1, S_2 , hence the same holds for the dimension of the corresponding polytopes (see Lemma 4.2). The following observation justifies the idea of decomposing slack matrices via k -sum.

Observation 4.13. *Let $k \geq 2$, and let $S \in \{0, 1\}^{m \times n}$ be a k -sum with factors S_1, S_2 , where S, S_1, S_2 are non-redundant slack matrices of polytopes P, P_1, P_2 respectively. Then $\text{rk}(S) > \max\{k, \text{rk}(S_1), \text{rk}(S_2)\}$. In particular, $\dim(P) > \max\{k - 1, \dim(P_1), \dim(P_2)\}$.*

Proof. Let $S = (S_1, (x_1, \dots, x_{k-1})) \oplus_k (S_2, (y_1, \dots, y_{k-1}))$ for some special rows $x_1, \dots, x_{k-1}, y_1, \dots, y_{k-1}$, such that none of S_1, S_2 is (isomorphic to) the identity matrix I_k . Let P, P_1, P_2 be the polytopes with slack matrices S, S_1, S_2 respectively. As a first remark, we have that $\text{rk}(S_1), \text{rk}(S_2) \geq k$: indeed, S_1 is an upper triangular block matrix with k blocks, the first $k-1$ consisting of a row of 1s, the last one being $S_1^{0 \dots 0}$, which has rank at least 1 (since S_1 does not have 0 columns), and the same holds for S_2 .

We now show that $\text{rk}(S) > \text{rk}(S_1)$, the proof for S_2 being the same. This will complete the proof. First, notice that $\text{rk}(S) \geq \text{rk}(S_1)$ since S_1 is a submatrix of S . Now, assume by contradiction that equality holds, hence there are $t = \text{rk}(S_1) = \text{rk}(S)$ columns of S_1 that form a basis for the column space of S_1 , and t corresponding columns of S that form a basis B for the column space of S . But then every column of S can be written in a unique way as linear combination of columns in B , implying that no two columns of S are identical when restricted to rows of S_1 , but different otherwise. Hence $S_2^{a_1 \dots a_{k-1}}$ consists of one column only for any $(a_1, \dots, a_{k-1}) \in \{(0, \dots, 0), (1, \dots, 0), \dots, (0, \dots, 1)\}$, in particular S_2 has exactly k columns. But as $\text{rk}(S_2) \geq k$, we conclude by Lemma 4.2 that P_2 has dimension $k-1$ and k vertices, i.e. P_2 is a simplex and $S_2 = I_k$, a contradiction. \square

4.3 Slack matrices of 2-level stable set polytopes

Stable set polytopes of perfect graphs are a prominent example of 2-level polytopes and it is natural to ask for an interpretation of the k -sum operation in this case. In this section we give a simple answer to this question, as well as describing an algorithm to recognize slack matrices of stable set polytopes in polynomial time. We recall (see Chapter 2, Section 2.3.1) that for a perfect graph $G(V, E)$, $\text{STAB}(G) = \{x \in \mathbb{R}_+^V : x(C) \leq 1 \text{ for all maximal cliques } C \text{ of } G\}$, and this description is non-redundant. Hence a 0/1, non-redundant slack matrix of $\text{STAB}(G)$ will have a column for each stable set of G and a row for each non-negativity inequality and for each maximal clique of G (we call such rows and inequalities clique rows and clique inequalities). In this section, we give a simple characterization of slack matrices of 2-level stable set polytopes that are k -sums. The characterization is based on the idea of composing two graphs G_1, G_2 by identifying the vertices of a clique of G_1 and a clique of G_2 , to obtain a graph G whose stable sets have a simple description in terms of the stable sets of G_1, G_2 . This goes back to [18] (see Theorem 4.1) and was studied in a more general setting in [22]. We say that a graph $G(V, E)$ has a clique cut-set K if V can be partitioned in V_1, V_2, K such that $V_1, V_2 \neq \emptyset, K$ is a clique, and there is no edge between V_1, V_2 . For simplicity, we allow K to be empty: this is equivalent to G being disconnected, in particular G is the disjoint union of G_1, G_2 with vertex sets V_1, V_2 respectively. It is easy to see that in this case $\text{STAB}(G) = \text{STAB}(G_1) \times \text{STAB}(G_2)$, hence if S, S_i denote the slack matrices of $\text{STAB}(G), \text{STAB}(G_i)$ for $i = 1, 2$, we have $S = S_1 \oplus S_2$. In general, if G has a clique cut-set of size $k-1$, and G_1, G_2 are the restriction of G to $V_1 \cup K, V_2 \cup K$ respectively, it is not hard to see that S is the k -sum of S_1, S_2 , with special rows corresponding to the non-negativity inequalities $x_v \geq 0$, for $V \in K$. We now show that some kind of converse

holds. We will assume, without loss of generality, that our slack matrices do not have constant (all zeros or all ones) rows.

Theorem 4.14. *Let $G(V, E)$ be a perfect graph and S be a 0/1 slack matrix of $\text{STAB}(G)$ with no constant rows, and let $k \geq 1$. Then S is a k -sum with special rows corresponding to non-negativity inequalities if and only if G has a clique cut-set of size $k - 1$.*

Proof. One direction has already been sketched above, so we focus on the other one. Assume that $S = (S_1, x_1, \dots, x_{k-1}) \oplus_k (S_2, y_1, \dots, y_{k-1})$ for some matrices S_1, S_2 and (if $k > 1$) special rows (x_1, \dots, x_{k-1}) of S_1 , and (y_1, \dots, y_{k-1}) of S_2 . We will prove that G has a clique cut-set of size $k - 1$ by induction on k .

- For $k = 1$, we have that $S = S_1 \oplus S_2$ and the non-negativity inequalities $z_v \geq 0$ for $v \in V$ belong either to S_1 or to S_2 , inducing a partition V_1, V_2 of V . As the slack of all the clique inequalities is determined by the slack of the non-negativity inequalities, it is easy to see that if one of V_1, V_2 was empty we would have that S_1 or S_2 consists of a single column, but then S has a constant row, a contradiction. Hence we have that $I \subset V$ is a stable set of G if and only if $I = I_1 \cup I_2$, where I_i is a stable set in V_i for $i = 1, 2$. This implies that there is no edge between V_1, V_2 , and we are done.
- Let $k \geq 2$. We have that the special rows correspond to non-negativity inequalities $z_{v_i} \geq 0$, for $i = 1, \dots, k - 1$, and since no column contains two 1's in correspondence of two special rows, the v_i 's form a clique K . Fix any special row, say $z_{v_1} \geq 0$: it is easy to see that the submatrix S' induced by the zeros of this row is the slack matrix of $\text{STAB}(G - v_1)$. But S' (after possibly removing a constant row) is a $k - 1$ -sum with respect to the rows $z_{v_i} \geq 0$ for $i = 2, \dots, k - 2$. Hence, by induction $G - v_1$ has $\{v_2, \dots, v_{k-1}\}$ as a clique cut-set. But then $\{v_1, \dots, v_{k-1}\}$ form a clique cut-set for G , and we are done.

□

Theorem 4.14 could be used to decompose the slack matrix of $\text{STAB}(G)$ into slack matrices of stable set polytopes of smaller graphs, for instance with the purpose of recognition, provided that G has a clique cut-set. However, we now describe a polynomial algorithm to recognize slack matrices of stable set polytopes of perfect graphs without using the notion of k -sums. First, as argued in Section 4.1.1, we can restrict ourselves to the case of non-redundant slack matrices. Given $S \in \{0, 1\}^{m \times n}$, let the rank of S be r , hence if S is a slack matrix of a polytope P the dimension of P is $d = r - 1$ thanks to Lemma 4.2. We know that if S is the slack matrix of $\text{STAB}(G)$ (we refer to this as the YES case), where G has d vertices, then there is a column of S (corresponding to the empty set) with exactly d zeros. Hence, if no such column exists, we output NO. Otherwise, for each such column c , we assume that the d rows that have zeros in position c correspond to the non-negativity inequalities. This allows us to reconstruct a graph G by connecting two vertices if they do not form a stable set, i.e. if

the two corresponding rows do not have a 1 in the same position. First, perfectness of G can be checked in polynomial time (see [23]). Now, the columns of S give us a list of stable sets of G , and by identifying the singletons among them we can easily obtain a list of maximal cliques of G from the rows of S which are not non-negativity rows. Provided that the entries of S do not give rise to any inconsistency while computing such lists, we need to check that these lists are complete. Now, it is well known that the maximal cliques of a graph (hence the independent sets as well) can be enumerated in total polynomial time (i.e., in time polynomial in the size of the output), for instance using the Bron–Kerbosch algorithm ([12]). Hence we can efficiently check whether the lists are complete in time polynomial in their size, i.e. polynomial in the size of S . If all these checks are successful, then S is the slack matrix of $\text{STAB}(G)$ and we output YES. In the worst case, we need to iterate the above procedure over all columns of S (with exactly d zeros) and, if the check fails for each of them, we output NO. We proved the following:

Theorem 4.15. *Let $S \in \{0, 1\}^{m \times n}$. Then there is an algorithm that runs in polynomial time in m, n and determines whether S is the slack matrix of $\text{STAB}(G)$, for some graph G .*

4.4 Recognition algorithms

In this section, we study the problem of recognizing k -sums: given a 0/1 matrix S and an integer $k \geq 1$, we want to determine whether S is a k -sum, and find the factors S_1, S_2 (and the special rows, if $k > 1$) such that $S = (S_1, x_1, \dots, x_{k-1}) \oplus_k (S_2, y_1, \dots, y_{k-1})$. Since we allow the rows and columns of S to be permuted in any way, the problem is non-trivial. In the following we will describe some algorithms to solve this problem. Our purpose is to use such algorithms to recognize slack matrices: given S a candidate slack matrix, we would like to decompose S as a k -sum of smaller matrices, which are all slack matrices (or can be turned into slack matrices by adding a row) if and only if S is, thanks to Lemma 4.11.

The starting point of our approach is the following observation: if a matrix S is a 1-sum $S_1 \oplus S_2$ with respect to a partition R_1, R_2 of its row set, then a column of the form (a, b) , where a, b are vectors corresponding to R_1, R_2 respectively, is a column of S if and only if a is a column of $S|_{R_1}$ and b is a column of $S|_{R_2}$. Moreover, the number of occurrences of (a, b) depends exclusively on the number of occurrences of a in $S|_{R_1}$ and of b in $S|_{R_2}$, or equivalently in S_1, S_2 . Intuitively, given uniform probability distributions on the columns of S , the probability of picking (a, b) in S is just the product of the probabilities of picking a in $S|_{R_1}$ and b in $S|_{R_2}$, as the latter are independent events. We will formalize and exploit this intuition in the next section, obtaining an algorithm for recognizing 1-sums. We will then extend this algorithm in order to recognize k -sums.

We remark that, even though we focus on 0/1 matrices, the results of this section can be easily extended to matrices with real entries or entries in any set.

4.4.1 Recognizing 1-sums via submodular function minimization

We recall some notions from information theory. We refer to [24] for a more complete exposition. Let A be a discrete random variable. For simplicity we overload notation and write $a \in A$ for a in the range of A , and $\Pr(a) = \Pr(A = a)$.

Definition 4.16. Let A, B two discrete random variables.

1. The *entropy* of A is: $H(A) = -\sum_{a \in A} \Pr(a) \log \Pr(a)$.
2. The *joint entropy* of A, B is: $H(A, B) = -\sum_{a \in A, b \in B} \Pr(a, b) \log \Pr(a, b)$.
3. The *conditional entropy* of A given B is: $H(A|B) = \sum_{b \in B} \Pr(b) H(A|B = b)$.
4. The *mutual information* of A, B is:

$$I(A, B) = \sum_{a \in A, b \in B} \Pr(a, b) \log \frac{\Pr(a, b)}{\Pr(a) \cdot \Pr(b)}.$$

We will use the following facts, whose proof can be found in [24]:

Proposition 4.17. Let A, B two discrete random variables. Then

1. $H(A, B) = H(A) + H(B|A)$.
2. $I(A, B) = H(A) - H(A|B) = H(B) - H(B|A)$.
3. $H(A|B) \leq H(A)$, with equality if and only if A, B are independent.
4. $I(A, B) \geq 0$, with equality if and only if A, B are independent.

Joint entropy extends the notion of entropy to pairs, or more generally to sets, of random variables: indeed, a set of random variables $\{X_1, \dots, X_m\}$ can be seen as a random variable X whose distribution is the joint distribution of the X_i 's. This can be applied to mutual information in a similar way, thanks to Proposition 4.17, part 2. We will now prove that entropy and mutual information, when considered as functions of sets of random variables, are submodular. Recall that a function $g : 2^\Omega \rightarrow \mathbb{R}$, where Ω is a finite set, is said to be *submodular* if for every $X, Y \subseteq \Omega$, $g(X) + g(Y) \geq g(X \cup Y) + g(X \cap Y)$. We are going to use the following alternative definition, which can be seen to be equivalent: g is submodular if and only if for every $X \subseteq Y \subseteq \Omega$, and for every $z \notin Y$,

$$g(X \cup \{z\}) - g(X) \geq g(Y \cup \{z\}) - g(Y). \quad (4.3)$$

In our context, Ω is a finite set of discrete random variables. For simplicity of notation, in the following we use capital letters to denote sets of random variables, $X \subseteq \Omega$ and lower case letters to denote individual random variables $z \in \Omega$.

Proposition 4.18. *The entropy $H : 2^\Omega \rightarrow \mathbb{R}_+$ is submodular.*

Proof. For every $X \subseteq Y \subseteq \Omega$ and $z \notin Y$, we claim:

$$H(X, z) - H(X) \geq H(Y, z) - H(Y).$$

Using Proposition 4.17, part 1, the left hand side is equal to $H(z | X)$ and the right hand side to $H(z | Y)$. At this point notice that conditioning cannot increase entropy (Proposition 4.17, part 3), thus:

$$H(z | Y) = H(z | X \cup (Y \setminus X)) \leq H(z | X).$$

The claim follows. □

Submodularity of entropy yields a similar property for mutual information.

Proposition 4.19. *The mutual information $I(X; \Omega \setminus X)$ is submodular as a function of $X \subseteq \Omega$.*

Proof. Fix $X \subseteq Y \subseteq \Omega$, $z \notin Y$. From Proposition 4.17, parts 1,2 it follows that

$$I(X; \Omega \setminus X) = H(X) + H(\Omega \setminus X) - H(\Omega).$$

Hence, we have that

$$\begin{aligned} A_X &:= I(X \cup \{z\}; \Omega \setminus (X \cup \{z\})) - I(X; \Omega \setminus X) \\ &= H(X \cup \{z\}) + H(\Omega \setminus (X \cup \{z\})) - H(X) - H(\Omega \setminus X) \\ &= H(X \cup \{z\}) - H(X) + H(\Omega \setminus (X \cup \{z\})) - H(\Omega \setminus X). \end{aligned} \tag{4.4}$$

We define A_Y analogously, thus it satisfies similar chain of identities as (4.4). By submodularity of entropy applied to $X \subseteq Y$ and to $\Omega \setminus (Y \cup \{z\}) \subseteq \Omega \setminus (X \cup \{z\})$, we conclude that $A_X \geq A_Y$, as desired. □

Let $S \in \{0, 1\}^{m \times n}$, and $X \subseteq [m]$ be a non-empty subset of row indices of S , and $\bar{X} = [m] \setminus X$. Consider the random variable C , that has uniform distribution over $\text{col}(S)$, i.e. it takes value c^* with probability equal $\frac{\mu(c^*)}{n}$, where $\mu(c^*)$ is the number of occurrences of the column c^* in $\text{col}(S)$. Let C_X be the restriction of C to the indices of X . Now, let $f : 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$ be defined as

$$f(X) = I(C_X, C_{\bar{X}}). \tag{4.5}$$

We remark that the definition of f depends on S , which we consider fixed throughout the section. Then f is non-negative and symmetric ($f(X) = f(\bar{X})$). Moreover, by identifying $\Omega = \{C_{\{i\}}, i \in [m]\}$ with $[m]$, we have that $f(X) = I(X, \Omega \setminus X)$ is submodular thanks to Proposition 4.19.

The next lemma shows that we can determine whether S is a 1-sum by minimizing f .

Lemma 4.20. *Let $S \in \{0, 1\}^{m \times n}$, and $\emptyset \neq X \subseteq [m]$ be a proper subset of row indices of S . Then S is a 1-sum with respect to X, \bar{X} if and only if $f(X) = 0$.*

Proof. Recall that $f(X) = 0$ if and only if C_X and $C_{\bar{X}}$ are independent random variables (Proposition 4.17, part 4). First, we prove ‘ \Rightarrow ’. Let S be a 1-sum with respect to X, \bar{X} , with corresponding decomposition $S = S_1 \oplus S_2$, where $S_i \in \{0, 1\}^{m_i \times n_i}$. Up to column permutation, we have that $S|_X$ consists of S_1 , repeated n_2 times, and $S_{\bar{X}}$ consists of S_2 , repeated n_1 times. Hence for any column $c = (c_X, c_{\bar{X}})$ of S , we have

$$\mu(c) = \mu_1(c_X)\mu_2(c_{\bar{X}}) = \frac{\mu(c_X)}{n_2} \frac{\mu(c_{\bar{X}})}{n_1},$$

where μ_i denotes the multiplicity of a column in S_i , $i = 1, 2$. Hence

$$\Pr(C = c) = \Pr(C_X = c_X, C_{\bar{X}} = c_{\bar{X}}) = \frac{\mu(c_X)}{n_2} \frac{\mu(c_{\bar{X}})}{n_1} \frac{1}{n} = \Pr(C_X = c_X) \Pr(C_{\bar{X}} = c_{\bar{X}}),$$

where we used $n = n_1 n_2$. This proves that C_X and $C_{\bar{X}}$ are independent.

We now prove ‘ \Leftarrow ’. Let a_1, \dots, a_h denote the different columns of $S|_X$, and b_1, \dots, b_k denote the different columns of $S|_{\bar{X}}$. Since C_X and $C_{\bar{X}}$ are independent, we have that, for any column $c = (a_i, b_j)$ of S ,

$$\mu(a_i, b_j) = n \cdot \Pr(C_X = a_i, C_{\bar{X}} = b_j) = n \cdot \Pr(C_X = a_i) \Pr(C_{\bar{X}} = b_j) = \frac{\mu_X(a_i) \mu_{\bar{X}}(b_j)}{n},$$

where $\mu_X(c), \mu_{\bar{X}}(c)$ denote the multiplicity of a column c of $S|_X, S|_{\bar{X}}$ respectively. Hence, if M denotes the matrix such that $M_{i,j} = \mu(a_i, b_j)$, we have that M is a matrix with nonnegative integer entries that has a rank 1 factorization of the form uv^\top , where $u_i = \mu_X(a_i)/n$, $v_j = \mu_{\bar{X}}(b_j)/n$, for $i = 1, \dots, h$, $j = 1, \dots, k$. Now, it is easy to see that one can turn this factorization into an integer one: let $u_1 = p_1/q_1$, where p_1, q_1 are coprimes, then q_1 must divide v_j for any j , since $u_1 v_j$ is integer. Then the factorization $q_1 u \cdot \frac{1}{q_1} v^\top = u'(v')^\top$ is such that v' is integer and u' has at least one more integer entry than u . Iterating, we obtain that $M = \bar{u} \bar{v}^\top$ where \bar{u}, \bar{v} have nonnegative integer entries. Now, let S_1 be the matrix consisting of the column a_i repeated \bar{u}_i times, for $i = 1, \dots, h$, and construct S_2 from \bar{v} in an analogous way. Then it is immediate to see that $S = S_1 \oplus S_2$ and in particular S is a 1-sum with respect to the row partition X, \bar{X} , which concludes the proof. \square

Notice that the previous proof also gives a way to efficiently reconstruct S_1, S_2 once we identified X such that $f(X) = 0$. In particular, if the columns of S are all distinct, then reconstructing S_1, S_2 is immediate: S_1 consists of all the distinct columns of $S|_X$, each taken once, and S_2 is obtained analogously from $S_{\bar{X}}$. The last ingredient we need is that every (symmetric) submodular function can be minimized in polynomial time:

Theorem 4.21. [83] *Let $f : 2^A \rightarrow \mathbb{R}$ be a symmetric submodular function. Then there is an algorithm that outputs a set X such that $X \neq \emptyset, A$ and $f(X)$ is minimum, using $O(n^3)$ calls to an oracle for f and $O(n^3)$ other basic operations, where $n = |A|$.*

As a consequence, we obtain the following:

Theorem 4.22. *Let $S \in \{0, 1\}^{m \times n}$. There is an algorithm that is polynomial in n, m and determines whether S is a 1-sum and, in case it is, outputs two matrices S_1, S_2 such that $S = S_1 \oplus S_2$.*

Proof. It is clear that $f(X)$ can be computed in polynomial time for any X . It suffices then to run Queyranne's algorithm to find X minimizing f . If $f(X) > 0$, then S is not a 1-sum. Otherwise, $f(X) = 0$ and S_1, S_2 can be reconstructed as described in the proof of Lemma 4.20. \square

We conclude the section with a decomposition result that will be useful for later. In light of Lemma 4.6, this result generalizes the fact that a polytope can be uniquely decomposed as cartesian product of "irreducible" polytopes (see [40] for a proof, in the context of abstract polytopes).

Lemma 4.23. *Let $S \in \{0, 1\}^{m \times n}$ be a 1-sum. Then there exists a partition $\{X_1, \dots, X_t\}$ of $[m]$ such that:*

1. *S is a 1-sum with respect to $X_i, \overline{X_i}$ for $i = 1, \dots, t$;*
2. *for any i and any X proper subset of X_i , S is not a 1-sum with respect to X, \overline{X} , i.e. the X_i 's are "minimal";*
3. *the partition X_1, \dots, X_t is unique up to permuting the labels.*

In particular, if S has all distinct columns, then there are matrices S_1, \dots, S_t such that $S = S_1 \oplus \dots \oplus S_t$, each S_i is irreducible, and the choice of the S_i 's is unique up to renaming and permuting columns.

Proof. We first recall the following well known property of submodular functions, which follows immediately from the definition: if a submodular function is minimized by X_1, X_2 , then it is also minimized by $X_1 \cap X_2$ (and $X_1 \cup X_2$). This means that the family of solutions of f of value 0 (which is non-empty thanks to Lemma 4.20, hence it has at least two elements for the symmetry of f) is closed under intersection: let X_1, \dots, X_t be the minimal (non-empty) sets of such family. Clearly, X_1, \dots, X_t are all disjoint. Moreover, since f is symmetric, the union of the X_i 's is $[m]$, hence X_1, \dots, X_t forms a partition of $[m]$ satisfying (i) (due to Lemma 4.20) and (ii) due to their minimality.

Moreover it easily follows that X_1, \dots, X_t are unique, as we now show: let X'_1, \dots, X'_t another partition of $[m]$ satisfying properties (i),(ii). If the two partitions are actually different, and not a renaming of each other, there exist X_i, X'_j with $X_i \neq X'_j$ and $X_i \cap X'_j \neq \emptyset$. But then $X_i \cap X'_j$ is an optimal solution which is a proper subset of (at least one of) X_i, X'_j , a contradiction.

To conclude, assume that S has all distinct columns. Then as argued above each S_i is obtained by picking each distinct column of $S|_{X_i}$ exactly once, and it is thus unique up to permutations, once X_i is fixed. Each S_i is irreducible thanks to the minimality of X_i and to Lemma 4.20. The fact that the X_i 's are unique up to renaming concludes the proof. \square

Remark 4.24. One can strengthen Lemma 4.23 in the following sense: let us call a matrix *repetitive* if it is formed by horizontally concatenating the same matrix multiple times, i.e. it is of the form $S = \begin{bmatrix} S' & \dots & S' \end{bmatrix}$ for some S' , or if it is isomorphic to a matrix of this form. Then repetitive matrices are the only ones that can admit multiple 1-sum factorizations, as in the following example:

$$\begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix} = \begin{bmatrix} 0 & 0 \end{bmatrix} \oplus \begin{bmatrix} 1 \end{bmatrix} = \begin{bmatrix} 0 \end{bmatrix} \oplus \begin{bmatrix} 1 & 1 \end{bmatrix}.$$

Formally, if S is a 1-sum and it is not repetitive, it has a factorization $S = S_1 \oplus \dots \oplus S_t$ in irreducible matrices that is unique up to renaming and permuting columns. We actually prove something stronger: assume that $S = S_1 \oplus S_2 = S'_1 \oplus S'_2$ where both 1-sums are with respect to the same partition R_1, R_2 , and $S_1 \neq S'_1$. Then S is repetitive. To see this, consider the matrix M defined in the proof of Lemma 4.20: M has a row for each column of $S|_{R_1}$, a column for each column of $S|_{R_2}$, and the entries are given by the multiplicity of each column in S . Then there are vectors u, v, u', v' with positive integers as entries and such that $M = uv^\top = u'(v')^\top$ and $u \neq u', v \neq v'$. But then we must have (without loss of generality) $u = \lambda u', v = \frac{1}{\lambda} v'$ for some integer $\lambda > 1$. This implies that all entries of M are divisible by λ , which implies in particular that S is repetitive.

4.4.2 Extension to k -sums

We now extend the previous results to obtain a polynomial algorithm to recognize k -sums, for constant $k > 1$. Recall that, if a 0/1 matrix S is a k -sum, then it has $k - 1$ special rows that divide S in submatrices $S^{0\dots 0}, S^{0\dots 1}, \dots, S^{1\dots 0}$, all of which are 1-sums *with respect to the same partition*. Hence, our algorithm starts by guessing the $k - 1$ special rows, and obtaining the corresponding submatrices $S^{0\dots 0}, S^{0\dots 1}, \dots, S^{1\dots 0}$. Let f_0, f_1, \dots, f_{k-1} denote the functions f as defined in (4.5) with respect to the matrices $S^{0\dots 0}, S^{0\dots 1}, \dots, S^{1\dots 0}$ respectively, and let $\tilde{f} = \sum_{i=0}^{k-1} f_i$. Notice that \tilde{f} is submodular, and is zero if and only if each f_i is. Let X be a proper subset of the non-special rows of S (which are the rows of any of $S^{0\dots 0}, S^{0\dots 1}, \dots, S^{1\dots 0}$). It is an easy consequence of Lemma 4.20 that $S^{0\dots 0}, S^{0\dots 1}, \dots, S^{1\dots 0}$ are 1-sums with respect to X if and only if $\tilde{f}(X) = 0$. Then S is a k -sum with respect to the chosen special rows if and only if the minimum of \tilde{f} is zero. Alternatively, one could first repeatedly decompose each $S^{a_1 \dots a_{k-1}}$ and

obtain a minimal partition of its row set, as described in Lemma 4.23, and then check whether these k partitions are refinements of a single partition X, \bar{X} of the row set. This is a simple combinatorial problem that can be solved efficiently.

Once a feasible partition is found, S_1, S_2 can be reconstructed by first reconstructing all $S_1^{a_1 \dots a_{k-1}}$'s, $S_2^{a_1 \dots a_{k-1}}$'s and then concatenating them and adding the special rows. We obtained the following:

Theorem 4.25. *Let $S \in \{0, 1\}^{m \times n}$, and $k \in \mathbb{Z}$ be a positive constant. There is an algorithm that is polynomial in n, m and determines whether S is a k -sum and, in case it is, outputs two matrices S_1, S_2 and special rows x_1, \dots, x_{k-1} of S_1 , y_1, \dots, y_{k-1} of S_2 , such that $S = (S_1, x_1, \dots, x_{k-1}) \oplus_k (S_2, y_1, \dots, y_{k-1})$.*

In order to apply Theorem 4.25 to decompose slack matrices, we need to deal with a last issue: in the algorithm, it is fundamental to guess the special rows that partition the column set in 1-sums. However, in principle there might be a slack matrix that is obtained as k -sum of other slack matrices, but where one of the special rows is redundant (i.e. it corresponds to a face, and not a facet, of the polytope). Then deleting such row still gives a slack matrix, but we cannot recognize such matrix as k -sum any more using our algorithm. However, the next lemma ensures that this does not happen, as long as we assume that the special rows are not redundant in the factors of the k -sum.

Lemma 4.26. *Let $S \in \{0, 1\}^{m \times n}$ and let $S_i \in \{0, 1\}^{m_i \times n_i}$ for $i = 1, 2$ such that $S = (S_1, x_1, \dots, x_{k-1}) \oplus_k (S_2, y_1, \dots, y_{k-1})$ for some special rows x_1, \dots, x_{k-1} of S_1 , and y_1, \dots, y_{k-1} of S_2 . Assume that S_1, S_2, S are slack matrices, and that the rows $x_1, \dots, x_{k-1}, y_1, \dots, y_{k-1}$ are non-redundant for S_1, S_2 respectively. Then the corresponding special rows in S are non-redundant as well.*

Proof. Assume by contradiction that r is a special row of S which is redundant, hence there exists another row r' of S such that $r' \leq r$ (i.e. r' has a zero in correspondence of every zero of r). Without loss of generality r' corresponds to a row r_1 of S_1 , i.e. r' consists (up to permutation) of r_1 repeated n_2 times, and similarly r corresponds to a special row of S_1 , say x_1 . But then it is clear that $r_1 \leq x_1$, i.e. x_1 is redundant in S_1 . \square

4.4.3 Numerical experiments

So far, we showed that the operation of k -sum decomposes a 0/1 slack matrix into two 0/1 slack matrices, and we provided an algorithm to recognize k -sums and provide such decomposition. Now, a natural question is how relevant is the operation of k -sum in the context of 2-level polytopes, in other words how many of the (finitely many) 2-level polytopes of dimension d can be constructed from lower dimensional polytopes via this operation. This is a non-trivial problem, in particular because it is not clear a priori whether the decomposition as k -sum is unique for $k \geq 2$, hence how many different combinations of lower dimensional

polytopes (and of the choices of the special rows) can give rise to the same polytope up to isomorphism. As a partial answer to this problem, we performed some numerical experiments on a database of 2-level polytopes of dimension up to 7, obtained by the authors of [10]¹. We implemented the algorithms from Theorem 4.22 and 4.25 in Matlab² and run them on the database. The results are listed in Table 4.1. As the table shows, there is some evidence that many 2-level polytopes (maybe a constant fraction of the total) are obtained via k -sum from lower dimensional 2-level polytopes, with consequences on their structure and extension complexity (see Corollary 4.12). Further research is needed to derive general estimates of the number of 2-level polytopes that are k -sums, and maybe to find other operations to represent 2-level polytopes that are not.

Dimension	1	2	3	4	5	6	7
1-sums	0	1	2	6	22	126	1276
2-sums			1	4	29	307	6435
3-sums				1	13	179	4786
4-sums					0	12	439
5-sums						0	7
6-sums							0
k -sums	0	1	3	11	64	624	12943
Total	1	2	5	19	106	1150	27292

Table 4.1 – The table lists the number of 2-level polytopes of a given dimension, from 1 to 7, that are k -sums, for $k = 1, \dots, 6$. The second-to-last row indicates the total number of k -sums, and the last row the total number of 2-level polytopes of each dimension.

4.4.4 Matroid polytopes

We now argue that the results in Section 4.4 imply that we can recognize the slack matrix of a 2-level base matroid polytope in polynomial time.

We will use basic notions of matroid theory that have already been defined in Chapter 2, Section 2.4. We also recall that, given a matroid $M(E, \mathcal{B})$, the *dual matroid* of M , denoted by M^* , is the matroid on E whose bases are the complements of the bases of M . We remark that, for any matroid M , the base polytopes $B(M)$ and $B(M^*)$ are affinely equivalent via the transformation $f(x) = 1 - x$ and hence have the same slack matrix.

We will use again Theorem 2.21 from [48]: the base polytope of a matroid M is 2-level if and only if M can be obtained from uniform matroids through direct sums and 2-sums. As the reader might imagine, the operation of direct sum of matroids is equivalent to the 1-sum of slack matrices of the corresponding base polytopes, and we will show a similar relation for the

¹The database can be found in <https://github.com/ulb/tl> and it has been extended to dimension 8. However, for reasons of computational power we only investigated polytopes of dimension at most 7.

²The code can be found at <http://disopt.epfl.ch/files/content/sites/disopt/files/users/249959/ksum.zip>

Chapter 4. Slack matrices of 2-level polytopes: recognition and decomposition

2-sum, even though we will need further assumptions (see 4.28 and 4.30 below). We remark that this correspondence breaks down for $k \geq 3$: the operation of 3-sum on matroids (whose definition can be found in [80]) does not correspond to the 3-sum of matrices. The general idea is to decompose our candidate slack matrix as 1-sum and 2-sum until each factor corresponds to the slack matrix of a uniform matroid. The latter can be easily recognized: indeed, as already remarked, the base polytope of the uniform matroid $U_{n,k}$ is the (n, k) -hypersimplex, described as

$$B(U_{n,k}) = \left\{ x \in \mathbb{R}^E : 0 \leq x_e \leq 1, \sum_e x_e = k \right\}.$$

If $2 \leq k \leq n-2$, the (non-redundant) slack matrix S of $B(U_{n,k})$ has $2n = 2|E|$ rows and $\binom{n}{k}$ columns of the form $\begin{bmatrix} \nu & \mathbf{1} - \nu \end{bmatrix}^T$ where $\nu \in \{0, 1\}^n$ is a vector with exactly k ones, hence can be recognized in polynomial time (in the size of S). We denote such matrix by $S_{n,k}$. If $k = 1$, or equivalently $k = n-1$, $S = S_{n,1} = S_{n,n-1}$ is just the identity matrix I_n . The case $k = 0$ or $k = n$ corresponds to a non-connected matroid whose base polytope is just a single vertex, and can be ignored for our purposes.

We will now investigate the relation between 1-sum and 2-sum of matroids and of the slack matrices of the corresponding base polytope.

We first need some preliminary assumptions. Let $M(E, \mathcal{B})$ be a matroid such that $B(M)$ is 2-level, and let S be a 0/1 slack matrix of $B(M)$. From now on we assume that:

1. M does not have loops or coloops.
2. S has a row for each inequality of the form $x(e) \geq 0$ for $e \in M$ (we refer to such rows as non-negativity rows).
3. S does not have any constant row (i.e. all zeros or all ones).

Assumption 1 is without loss of generality as, if e is a loop or coloop of M , then $B(M)$ has a constant coordinate in correspondence of e and is thus isomorphic to $B(M - e)$.

We now justify the Assumption 2. In general, the non-negativity inequalities are not necessarily facet-defining. However, we claim that we can always assume that they correspond to rows of S , and in particular this does not affect whether S is a 1-sum or a 2-sum. Let us first consider the case in which M is connected. The following is a well known result that can be found for instance in [95]:

Lemma 4.27. *Let M be a connected matroid, then for any $e \in E$, at least one of $M - e$, M/e is connected.*

From Theorem 2.28 it follows that the inequality $x(e) \geq 0$ (resp. $x(e) \leq 1$) is facet defining for $B(M)$ if and only if $M - e$ (resp. M/e) is connected. Then by Lemma 4.27 one of the two

inequalities is always facet-defining for any $e \in E$, i.e. one of the two corresponding rows r is always in S : but the other row is equal to $\mathbf{1} - r$, it hence clearly can always be added to S : $S + (\mathbf{1} - r)$ is a 1-sum (resp. 2-sum) if and only if S is. Now, if M is not connected, then for any $e \in E$ there is a connected matroid M_1 with $M = M_1 \times M_2$ and e element of M_1 . Moreover, it is easy to see that $B(M) = B(M_1) \times B(M_2)$, and that $x(e) \geq 0$ (or $x(e) \leq 1$) is facet defining for $B(M)$ if and only if it is for $B(M_1)$. This proves our claim.

Finally, as remarked in Section 4.1.1, Assumption 3 is without loss of generality, and, thanks to Assumption 1, it is not in contradiction with Assumption 2 (clearly a non-negativity row is constant if and only if the corresponding element is a loop or a coloop).

Now, let us focus on the operation of 1-sum. On one hand, it is clear that, if S_1, S_2 are the slack matrices of $B(M_1), B(M_2)$ respectively, then $S_1 \oplus S_2$ is the slack matrix of $B(M_1) \times B(M_2) = B(M_1 \oplus M_2)$. We now need to prove that the converse holds, i.e. we need to make sure that, whenever we decompose the slack matrix of a matroid base polytope as a 1-sum, the factors are still matroid base polytopes.

Lemma 4.28. *Let $M(E, \mathcal{B})$ be a matroid such that $B(M)$ admits a slack matrix S with 0/1 entries. If $S = S_1 \oplus S_2$ for some matrices S_1, S_2 , then there are matroids M_1, M_2 such that $M = M_1 \oplus M_2$ and S_i is a slack matrix of $B(M_i)$ for $i = 1, 2$.*

Proof. By Assumption 2, S contains all the rows corresponding to inequalities $x(e) \geq 0$, for any e element of M . Each such non-negativity inequality belongs either to S_1 or to S_2 , hence we can partition E into E_1, E_2 accordingly. Notice that none of E_1, E_2 can be empty: if for instance E_2 is empty, then all the rows corresponding to $x(e) \geq 0$ belong to the partition R_1 (defined as usual). But then the slack of a vertex with respect to every other inequality (of form $x(U) \leq \text{rk}(U)$) depends entirely on the slack with respect to the rows in R_1 , implying that a column of $S|_{R_1}$ can be completed to a column of S in a unique way. Hence, since S is a 1-sum, we must conclude that S_2 is made of a single column, contradicting the fact that S does not have constant rows (Assumption 3).

Now, let $\mathcal{B}_i = \{B \cap E_i : B \in \mathcal{B}\}$ for $i = 1, 2$. By definition of 1-sum of matrices, $B(M) = \{B_1 \cup B_2 : B_i \in \mathcal{B}_i \text{ for } i = 1, 2\}$. This implies that $M = M_1 \oplus M_2$ where $M_i = M|_{E_i}$ for $i = 1, 2$, thus $B(M) = B(M_1) \times B(M_2)$. Hence, for every row of S corresponding to an inequality $x(U) \leq \text{rk}(U)$, we have either $U \subset E_1$, $U \subset E_2$, or the inequality is redundant and can be removed. In the first case, clearly the row is in R_1 as its entries depend only on the rows $x(e) \geq 0$ for $e \in E_1$, and similarly in the second case the row is in R_2 . As by removing redundant rows we do not change the polytopes of which S, S_1, S_2 are slack matrices, we then conclude that S_i is a slack matrix of $B(M_i)$ for $i = 1, 2$. \square

Corollary 4.29. *Let M be a matroid such that $B(M)$ has a 0/1 slack matrix S . Then M is connected if and only if S is not a 1-sum.*

We now deal with slack matrices of connected matroids and with the operation of 2-sum. We already dealt with the case of uniform matroids. To consider matroids which are 2-sums, we recall that, by Theorem 2.21, every connected matroid whose base polytope is 2-level is a 2-sum of uniform matroids, which form the vertices of a tree as in Theorem 2.33. Hence, by picking a leaf of the tree, we can assume that our candidate slack matrix is 2-sum of a uniform matroid and some other matroid.

Lemma 4.30. *Let $M(E, \mathcal{B})$ be a connected matroid such that $B(M)$ has a 0/1 slack matrix S . Assume there are S_1, S_2 such that $S = (S_1, x_1) \oplus_2 (S_2, y_1)$, let $S'_1 = S_1 + (\mathbf{1} - x_1)$ and similarly for S'_2 . Assume that S_1 or S'_1 is equal to $S_{n,k}$ for some $n > k \geq 1$. Then there is a matroid M_2 such that $M = U_{n,k} \oplus_2 M_2$ and S'_2 is a slack matrix of $B(M_2)$.*

Proof. As already argued we can assume that S has rows corresponding to all the non-negativity inequalities. We first claim that the special row r does not correspond to any non-negativity inequality: indeed, if it corresponds to $x(e) \geq 0$ for some $e \in E$, then it is not hard to see that S^{00} is the slack matrix of $M - e$, and similarly S^{11} is the slack matrix of M/e . But both matrices are 1-sums, hence by Corollary 4.29, none of $M - e, M/e$ is connected, a contradiction by Lemma 4.27. Hence, each inequality $x(e) \geq 0$ corresponds to a row of either S_1 or S_2 , giving a partition of E in E_1, E_2 . We will now proceed similarly as in the proof of Lemma 4.28: first, by noticing that the slack of any vertex with respect to $x(U) \leq \text{rk}(U)$ depends exclusively on the slack with respect to the non-negativity inequalities, we can again conclude that E_1, E_2 are not empty. Since $S_1 = S_{n,k}$ is the slack matrix of $U_{n,k}$, the special row x_1 of S_1 corresponds to the inequality $x(p) \geq 0$, or $x(p) \leq 1$ for some element p : we can assume that S_1 contains both rows (which are opposite), so that we do not need to mention S'_1 , and similarly for S_2 , and we consider the case in which x_1 corresponds to $x(p) \geq 0$, the other being analogous. Notice that p is not in E , as the special row of S does not correspond to a non-negativity inequality. Let us define $M_1 = U_{n,k}$ on ground set $E'_1 = E_1 + p$, with base set $\mathcal{B}_1 = \binom{E'_1}{n}$ and let:

$$\mathcal{B}_2 = \{B_2 + p : B_1 \cup B_2 \in \mathcal{B}, B_1 \subset E_1, |B_1| = k\} \cup \{B_2 : B_1 \cup B_2 \in \mathcal{B}, B_1 \subset E_1, |B_1| = k-1\}.$$

We remark that the two sets forming \mathcal{B}_2 are both non-empty and, due to the 2-sum structure of S , a set $B_2 \in \mathcal{B}_2$ can be completed to a basis of M by adding any $B_1 \in \mathcal{B}_1$ that satisfies $p \in B_1 \Delta B_2$, and removing p . Hence, if we show that M_2 with ground set $E'_2 = E_2 + p$ and base set \mathcal{B}_2 is a matroid, we will have that $M = M_1 \oplus_2 M_2$. In particular we now show that \mathcal{B}_2 satisfies the axioms for the base set of a matroid: it is non-empty (which is clear) and for any $B_2, B'_2 \in \mathcal{B}_2$ and $e \in B_2 \setminus B'_2$, there exists $f \in B'_2 \setminus B_2$ such that $B_2 - e + f \in \mathcal{B}_2$. We fix such B_2, B'_2, e and distinguish a number of cases.

1. $p \in B_2 \cap B'_2$. Then for any $B_1 \in \mathcal{B}_1$ with $p \notin B_1$ we have that $B_1 \cup B_2 - p, B_1 \cup B'_2 - p$ are bases of M , hence by applying the base axiom to them we obtain that there exists $f \in (B_1 \cup B'_2 - p) \setminus (B_1 \cup B_2 - p) = B'_2 \setminus B_2$ such that $B_1 \cup B_2 - p - e + f \in \mathcal{B}$, but then $B_2 - e + f \in \mathcal{B}_2$ by definition.

2. $p \notin B_2 \cup B'_2$. This case is analogous to the previous one.
3. $p \in B_2 \setminus B'_2$, and $e \neq p$. Let $B_1, B'_1 \in \mathcal{B}$ with $B_1 \Delta B'_1 = \{p, g\}$ and in particular $p \in B'_1 \setminus B_1$. Then we have $B = B_1 \cup B_2 - p, B' = B'_1 \cup B'_2 - p \in \mathcal{B}$. Then by the base axiom there exists $f \in B' \setminus B$ with $B - e + f \in \mathcal{B}$. Since $g \in B_1 \setminus B'_1 \subset B \setminus B'$, we have $g \neq f$ and we conclude that $f \in B'_2 \setminus B_2$, hence $B_2 - e + f \in \mathcal{B}_2$.
4. $p \in B_2 \setminus B'_2$, and $e = p$. This case is analogous to the previous one, but we apply the axiom to $g \in B \setminus B'$ instead of e .
5. $p \in B'_2 \setminus B_2$. Let $B_1, B'_1 \in \mathcal{B}$ with $B_1 \Delta B'_1 = \{p, g\}$ and in particular $p \in B_1 \setminus B'_1$, then again $B = B_1 \cup B_2 - p, B' = B'_1 \cup B'_2 - p \in \mathcal{B}$ and there is $f \in B' \setminus B$ with $B - e + f \in \mathcal{B}$. If $f \in B'_2$, then $f \notin B_2$ and we are done as before. Otherwise $f = g$, then $B - e + g = B'_1 \cup B_2 - e \in \mathcal{B}$, but then by definition $B_2 - e + p \in \mathcal{B}_2$.

Now, we have $M = M_1 \oplus_2 M_2$ by construction. Hence $B(M)$ is isomorphic to $B(M_1) \times B(M_2) \cap \{x_p + y_p = 1\}$, thanks to Lemma 2.24, and can hence be described by: a description of $B(M_1)$, a description of $B(M_2)$, and the equalities $x(E_1) + y(E_2) = \text{rk}(M)$, $x_p + y_p = 1$, which do not appear in the slack matrix. Now, as S_1 is the slack matrix of $B(M_1)$, the rows of R_2 must correspond to a description of $B(M_2)$: from this we conclude that S_2 is a slack matrix of $B(M_2)$. \square

We are now ready for the main theorem of the section.

Theorem 4.31. *Let $S \in \{0, 1\}^{n \times m}$. There is an algorithm that is polynomial in n, m and decides whether S is the slack matrix of $B(M)$ for a matroid M .*

Proof. First, we need to ensure that, if M_1, M_2 are connected, $M = M_1 \oplus_2 M_2$ and S_i is a slack matrix of $B(M_i)$ for $i = 1, 2$, then the slack matrix of $B(M)$ is actually the 2-sum of S_1, S_2 for some special rows x_1, y_1 . If p is the common element of M_1, M_2 , x_1 must be the row corresponding to $x(p) \geq 0$, and y_1 the row corresponding to $y(p) \leq 1$, or viceversa. Then, to conclude that S is equal to $(S_1, x_1) \oplus_2 (S_2, y_2)$, we only need the special rows to be non-redundant, as this ensures that no row is lost when doing the 2-sum (see Lemma 4.26). We already mentioned that the inequality $x(p) \geq 0$ (resp. $x(p) \leq 1$) is non-redundant for $B(M)$ if and only if $M - p$ (resp. M / p) is connected. Moreover, thanks to Theorem 2.21, we can consider each M_i as a 2-sum of uniform matroids, among which there will be $U_{n_1, k_1}, U_{n_2, k_2}$ containing the element p . Since our matroids are connected, by Proposition 2.23 we can assume that $n_i \geq 3$ and $1 \leq k_i \leq n_i - 1$ for $i = 1, 2$. It can be easily checked that, for any M', M'_1, M'_2 with $M' = M'_1 \oplus_2 M'_2$, and e element of M'_1 , $M' - e = (M'_1 - e) \oplus_2 M'_2$ and $M' / e = (M'_1 / e) \oplus_2 M'_2$; moreover, M' is connected if and only if M'_1, M'_2 are. Using this, we can focus on whether $U_{n_i, k_i} - p, U_{n_i, k_i} / p$ are connected for $i = 1, 2$. Notice that by Lemma 4.27 one of the two must be connected. If for one i we have $2 \leq k_i \leq n_i - 2$, then both of $U_{n_i, k_i} - p, U_{n_i, k_i} / p$ are connected, hence we can always find a suitable couple of non-redundant special rows. If, $k_1 = 1$, and $k_2 = n_2 - 1$, $U_{n_1, 1} - p = U_{n_1 - 1, 1}$ is connected (since $n_1 \geq 3$) and similarly $U_{n_2, n_2 - 1} / p = U_{n_2 - 1, n_2 - 2}$

is connected, hence we find our non-redundant special rows. The only problematic case arises when $k_1 = k_2 \in \{1, n_i - 1\}$ for $i = 1, 2$. But we can assume that the latter case is never verified as we have $U_{n_1,1} \oplus_2 U_{n_2,1} = U_{n_1+n_2-2,1}$ and similarly $U_{n_1,n_1-1} \oplus_2 U_{n_2,n_2-1} = U_{n_1+n_2-2,n_1+n-3}$.

Now, let S be a 0/1 matrix, we summarize the process of determining whether S is a slack matrix of $B(M)$ for some matroid M . First, we check whether $S = S_{n,k}$ for some n and k , in which case we are done. Then, we run the algorithm to recognize 1-sums, and if S is a 1-sum, we decompose it in factors S_1, \dots, S_t which are not 1-sums and test each S_i separately. This can be done efficiently thanks to Theorem 4.22, and using Lemma 4.28 we have that S is the slack matrix of $B(M)$ if and only if S_i is the slack matrix of $B(M_i)$ for each i , and $M = M_1 \oplus \dots \oplus M_t$. We can now assume that S is irreducible (i.e. it is not a 1-sum). In order to apply Lemma 4.30, we need to check whether S is a 2-sum with respect to any special row, until we detect a decomposition where one of the two factors has form $S_{n,k}$ (informally, such factor corresponds to a leaf in the tree decomposition of M). Then we continue on the other factor, until we decompose S as a repeated 2-sum of matrices S_1, \dots, S_t where $S_i = S_{n_i,k_i}$ for $i = 1, \dots, t$ (of course, if this is not possible, we conclude that S is not a slack matrix of a base polytope). The S_i 's form a tree structure T similarly as the factors of a 2-sum of a matroid, but their 2-sum does not necessarily correspond to a matroid. Indeed, each S_i is the slack matrix of both U_{n_i,k_i} and its dual, and the form of each special row (whether $x(p) \geq 0$ or $x(p) \leq 1$) must be coherent: if S_1 is the slack matrix of U_{n_1,k_1} , this determines the form of its special row, and of the special row of each neighbor of S_1 , but conflicts may arise as, if some of the S_i is the identity matrix, then the form of their special row is fixed. This problem can be seen as trying to color a tree with two colors, where some nodes can have a predetermined color. However, if there exists a feasible coloring, then this coloring determines a matroid M , and is essentially unique: it is easy to see that the only other possible coloring gives rise to the dual matroid M^* . To check whether there exist a feasible coloring can be done efficiently, and this concludes the algorithm. Notice that, in case S is the slack matrix of $B(M)$, M (or its dual) can be reconstructed by successively taking the 2-sum of the S_i 's. \square

4.5 Matroid polytopes: an alternative approach

In this section we describe an alternative approach to recognize slack matrices of 2-level base matroid polytopes, in the case the matroid is connected. This, together with Theorem 4.22 and Lemma 4.28, provides an alternative proof of Theorem 4.31 which is not based on 2-sum of slack matrices. This approach is based on some properties of (not necessarily 2-level) matroid polytopes which might be of independent interest, as they offer new connections between the facial structure of a matroid polytope and the structure of the matroid itself. Moreover, as most of the proof does not use 2-levelness, the results in this section could be extended to recognize slack matrices of general matroid polytopes (see Remark 4.39).

Throughout the section we assume that $M(E, \mathcal{B})$ is a connected matroid. Given $S \in \{0, 1\}^{m \times n}$, we want to decide whether it is a slack matrix of $B(M)$ for some M . If it is, we say we are

in the YES case. As remarked in Section 4.1.1, we can assume that, in the YES case, S is a non-redundant slack matrix.

The very first step of the algorithm is to obtain a graph H from S that, if S is a slack matrix of P , is the skeleton of P . We recall that the skeleton of P is the graph whose vertices are the vertices of P , and where two vertices are adjacent if they are in P (i.e. if their segment is a 1-dimensional face of P). It is easy to see that we can efficiently obtain H : for any two columns i, j , consider all the rows r with $S_{ri} = S_{rj} = 0$. Vertices i, j are adjacent if and only if there is no other column for which all of these rows have value 0. The structure of the skeleton of matroid polytopes is well known, see for instance [39]: two vertices χ^{B_1}, χ^{B_2} of $B(M)$ are adjacent if and only if $|B_1 \triangle B_2| = 1$. In this case we abuse notation and say that B_1, B_2 are adjacent, and write $B_1 = B_2 + e - f$ for some $e, f \in E$.

We will proceed as follows: we first assume that we are in the YES case and prove some properties of the skeleton H . In particular, we show that a maximal clique in H corresponds to a circuit or a cocircuit of M , and identify from this which rows of the slack matrix induce the inequalities $0 \leq x \leq 1$, i.e. what are the elements of the ground set of M . This naturally implies an efficient algorithm that, given S as input, produces a list \mathcal{L} that in the YES case is the set $\mathcal{C}(M)$ of the circuits of a matroid M . From that, we reconstruct M and the vertices and facets of its base polytope, hence we can compute the slack matrix of $B(M)$ and verify whether we are in the YES case or not.

Recall the non-redundant description of $B(M)$ given in Equation (2.6). We call the elements e such that M/e is connected contractible elements, and similarly the elements e such that $M - e$ is connected are called deletable. We remark that contractible elements are flacets, but here we want to consider them separately. A facet-defining inequality of the form $x(e) \geq 0$ or $x(e) \leq 1$ will be called an *element inequality* (and the corresponding row in the slack matrix will be an *element row*), and an inequality of the form $x(F) \leq \text{rk}(F)$ for $F \in \mathcal{F}$ with $|F| \geq 2$ will be called a *facet inequality* (and the corresponding row will be a *facet row*).

4.5.1 Phase 1: finding the circuits of M

In this section we observe a correspondence between the cliques of H , the skeleton of $B(M)$, and the circuits and cocircuits of M . It is easy to notice that a circuit C in M generates maximal cliques of size $|C|$ in H : each clique consists of bases that are equal everywhere except on C , and each basis lacks a different element of C . In what follows, we are going to focus on circuits of size at least 3. The reasons of this will be made clear at the end of the section. We remark that the results of this subsection do not use the fact that $B(M)$ is 2-level.

Lemma 4.32. *Let B_1, B_2, B_3 be three distinct bases of M such that $\Delta = \{B_1, B_2, B_3\}$ is a clique of H . Then there is a unique maximal clique $K = \{B_1, \dots, B_k\}$ containing Δ , and there are elements e_1, \dots, e_k such that one of the following holds:*

1. $e_i \in B_j$ if and only if $i \neq j$, and e_1, \dots, e_k is a circuit of M (we say that K induces a circuit

in M).

2. $e_i \in B_j$ if and only if $i = j$, and e_1, \dots, e_k is a cocircuit of M (we say that K induces a cocircuit in M).

Proof. Since B_1, B_2, B_3 are pairwise adjacent, we have that there exist $e_{12}, e_{21}, e_{13}, e_{31}, e_{23}, e_{32} \in E$ with $e_{12} \neq e_{21}, e_{13} \neq e_{31}, e_{23} \neq e_{32}$, such that

$$B_1 = B_2 - e_{12} + e_{21} = B_3 - e_{13} + e_{31}, \quad B_2 = B_3 - e_{23} + e_{32},$$

which implies

$$B_3 - e_{13} - e_{21} + e_{31} + e_{12} = B_3 - e_{23} + e_{32}.$$

We consider two cases:

1. $e_{13} = e_{12}$, which implies $e_{21} = e_{23}$ and $e_{31} = e_{32}$. We can then simplify the notation and write e_1 for e_{13}, e_{12} , e_2 for e_{21}, e_{23} , e_3 for e_{31}, e_{32} . The previous relations become:

$$B_1 = B_2 - e_1 + e_2 = B_3 - e_1 + e_3, \quad B_2 = B_3 - e_2 + e_3,$$

in other words we have $e_i \in B_j$ if and only if $i = j$ for $i, j \in [3]$. Now, we proceed to prove that there is a unique maximal clique containing Δ . Let B be a basis that is adjacent to B_1, B_2, B_3 (in short we write that B is adjacent to Δ). We first claim that $e_1, e_2, e_3 \in B$. Indeed, assume by contradiction that $e_1 \notin B$. Then $\{e_1\} = B_2 \setminus B$, hence $B = B_2 - e_1 + e$ for some $e \neq e_2$ (otherwise $B = B_1$). Hence $e_2 \notin B$, but then $\{e_1, e_2\} \subseteq B_3 \setminus B$, a contradiction to the fact that B, B_3 are adjacent. Now, consider two distinct B, B' bases adjacent to Δ , we show that B, B' must be adjacent. If they are not, there are elements $e, f \in E$ such that $\{e, f\} \subseteq B' \setminus B$. Since B, B_1 are adjacent, we can assume without loss of generality that $e \notin B_1$, but then $\{e\} = B' \setminus B_1 = \{e_1\}$, in contradiction with the fact that $e_1 \in B$ must hold. It then follows that the maximal clique $K = \{B_1, \dots, B_k\} \supseteq \Delta$ is unique. If $k \geq 4$, consider B_4 , which contains e_1, e_2, e_3 . We have $B_4 = B_1 - e + e_1 = B_2 - f + e_2 = B_3 - g + e_3$, but since $B_1 + e_1 = B_2 + e_2 = B_3 + e_3$, we get $e = f = g$ and we can denote this element by e_4 (note that it is clearly distinct from e_1, e_2, e_3). Iterating the argument for $B_i, i = 5, \dots, k$, we can obtain $e_5, \dots, e_k \in E$ and verify that for any $i, j \in [k]$, $e_i \in B_j$ if and only if $i = j$. Now, from the maximality of K it follows that $\{e_1, \dots, e_k\}$ is the unique circuit contained in $B_1 + e_1 = \dots = B_k + e_k$.

2. We now assume $e_{13} \neq e_{12}$. Now, if $e_{31} \neq e_{21}$, we must have $e_{13} = e_{21} = e_{23}$, $e_{31} = e_{12} = e_{32}$, but this implies $B_2 = B_3$, a contradiction. Hence we have $e_{31} = e_{21}$, that implies $e_{13} = e_{23}$, $e_{12} = e_{32}$. We will now proceed analogously as the previous case. We can write e_1 for e_{31}, e_{21} , e_2 for e_{12}, e_{32} , e_3 for e_{13}, e_{23} . The previous relations become:

$$B_1 = B_2 - e_2 + e_1 = B_3 - e_3 + e_1, \quad B_2 = B_3 - e_3 + e_2,$$

in other words we have $e_i \in B_j$ if and only if $i = j$ for $i, j \in [3]$. Let B be a basis that is adjacent to B_1, B_2, B_3 (in short we write that B is adjacent to Δ). We first claim that $e_1, e_2, e_3 \notin B$. Indeed, assume by contradiction that $e_1 \in B$. Then $B = B_2 - e + e_1$ for some $e \neq e_2$, otherwise $B = B_1$. But then $e_2 \in B$ hence $\{e_1, e_2\} \subseteq B \setminus B_3$, a contradiction. Now, consider two distinct B, B' bases adjacent to Δ , we show that B, B' must be adjacent. If they are not, there are elements $e, f \in E$ such that $\{e, f\} \subseteq B' \setminus B$. Now, since B', B_1 are adjacent we can assume without loss of generality that $e \in B_1$. But, since $e \neq e_1$, we have $\{e, e_1\} \subseteq B_1 \setminus B$, a contradiction with the fact that B, B_1 are adjacent. Hence we proved that in this case as well the maximal clique $K = \{B_1, \dots, B_k\} \supseteq \Delta$ is unique. If $k \geq 4$, consider B_4 , we have $e_1, e_2, e_3 \notin B_4$, hence $B_4 = B_1 - e_1 + e = B_2 - e_2 + f = B_3 - e_3 + g$ and we conclude $e = f = g := e_4$. Iterating the argument for B_i , $i = 5, \dots, k$, we can obtain $e_5, \dots, e_k \in E$ and verify that for any $i, j \in [k]$, $e_i \in B_j$ if and only if $i = j$. Now, let $D = \{e_1, \dots, e_k\}$, we need to show that D is a cocircuit of M , i.e. D is dependent in M^* (equivalently, $D \cap B \neq \emptyset$ for any $B \in \mathcal{B}$) and is minimal with this property. The minimality follows immediately as for any $i = 1, \dots, k$ one has $B_i \cap (D - e_i) = \emptyset$. To prove that D is dependent in M^* , assume by contradiction that there is $B \in \mathcal{B}$ with $D \cap B = \emptyset$. Then, applying the basis exchange axiom to B_1, B , we get that there is $e \in B \setminus B_1$ such that $B' = B_1 - e_1 + e \in \mathcal{B}$. Now, $B' \cap D = \emptyset$ and B' is adjacent to B_1 . Since K is maximal and $B \notin K$, without loss of generality B is not adjacent to B_2 . But $B' = B_1 - e_1 + e = B_2 - e_2 + e$, a contradiction.

□

To prove the next lemma, we need a few facts.

Observation 4.33. *Let M a matroid, C (respectively D) be a circuit (cocircuit) of M and $F \subset E$.*

1. *If F is a flat of M , $|C| \geq 2$ and $|C \cap F| \geq |C| - 1$, then $C \subset F$.*
2. *If $M|F$ has no coloops, then $E - F$ is a flat of M^* .*
3. *If $M|F$ has no coloops, $|D| \geq 2$, and $F \cap D \neq \emptyset$, then $|F \cap D| \geq 2$.*
4. *Let M be connected. Then F is a flacet of M if and only if $E - F$ is a flacet of M^* .*

Proof. 1. Assume that $C \cap F = C - e$ for some $e \in C$. $C - e$ is independent, hence it is contained in a basis B of $M|F$. Since F is a flat, we have $\text{rk}(F + e) > \text{rk}(F) = |B|$, but then in $M|(F + e)$ B is a independent set which cannot be extended to any basis, a contradiction.

2. $E - F$ is a flat of M^* if and only if for any $e \in F$, $\text{rk}^*(E - F + e) > \text{rk}^*(E - F)$, where rk^* is the rank function of M^* . Since $\text{rk}^*(A) = \text{rk}(E - A) + |A| - \text{rk}(E)$ for any $A \subset E$, the latter is equivalent to $\text{rk}(F - e) + 1 > \text{rk}(F)$, which holds since $M|F$ has no coloops.

3. This immediately follows from the previous two statements: we have that $E - F$ is a flat of M^* and D is a circuit of M^* , hence $|F \cap D| < 2$ would imply $|D \cap (E - F)| \geq |D| - 1 \implies D \subset E - F$, in contradiction with $F \cap D \neq \emptyset$.
4. The statement immediately follows from the fact that a matroid is connected if and only if its dual is, and from the following relations between contraction, restriction and deletion: $M|F = M \setminus (E - F) = (M^* / (E - F))^*$, and $M/F = (M^* | (E - F))^*$. These facts are well known and can be found for instance in [80].

□

Observation 4.34. *Let $M = M_1 \oplus_2 M_2$.*

1. *Let $e \in E_1 - p$ (resp. $E_2 - p$). Then $M - e = (M_1 - e) \oplus_2 M_2$ (resp. $M_1 \oplus_2 (M_2 - e)$).*
2. *Let $C \subset E$ such that $C \cap E(M_i) \neq \emptyset$ for $i = 1, 2$. Then C is a circuit of M if and only if $C_i = (C \cap E_i) + p$ is a circuit of M_i for $i = 1, 2$.*

Proof. We only prove the second fact as the first follows immediately from the definition of 2-sum. First notice that if I_1, I_2 are independent sets of M_1, M_2 respectively, such that $p \notin I_1 \cup I_2$ and $I_1 \cup I_2$ is dependent in M , then $I_i + p$ is dependent in M_i for $i = 1, 2$. Let C be a circuit of M . For $i = 1, 2$, $C \cap E_i$ is independent (in M and in M_i), and C is dependent in M , hence C_i is dependent in M_i . We are left to show that, for any $e \in C_i$, $C_i - e$ is independent in M_i . This is clear for $e = p$, so let $e \neq p$. We consider without loss of generality the case $i = 1$. We know that $C - e$ is independent in M , so $C - e \subset B_1 \cup B_2 - p$ for bases B_1 of M_1 and B_2 of M_2 with $p \in B_1 \triangle B_2$. Moreover, $C - e = C_1 \cup C_2 - p - e$. Hence $C_2 - p \subset B_2$ but since C_2 is dependent in M_2 , $p \notin B_2$, hence $p \in B_1$. Therefore $C_1 - e \subset B_1$ i.e. $C_1 - e$ is independent in M_1 .

Now let C_1, C_2 be circuits in M_1, M_2 . If C is independent in M , then either C_1 is independent in M_1 or C_2 is dependent in M_2 , a contradiction. Let $e \in C$, assume without loss of generality $e \in E_1$. Note that $C - e \subset C_1 - e \cup C_2 - p$. By definition, $C_1 - e$ is independent in M_1 and $C_2 - p$ is independent in M_2 . Since by definition of 2-sum p is not a loop of M_1 (resp. a coloop of M_2), we extend $C_1 - e$ (resp. $C_2 - p$) to a basis B_1 of M_1 (resp. B_2 of M_2) containing p (resp. not containing p). As $C - e \subseteq C_1 - e \cup C_2 - p \subseteq B_1 \cup B_2 \in \mathcal{B}(M)$, we conclude that $C - e$ is independent in M . □

Observation 4.35. *Let M be 3-connected such that $M - e$ is not connected. Then $M = U_{3,2}$.*

Proof. Recall that M is k -connected if it has no s -separation for $s = 1, \dots, k - 1$, where an s -separation is a set $X \subset E$ such that $|X|, |E - X| \geq s$, and $\text{rk}(X) + \text{rk}(E - X) < \text{rk}(M) + s$. Since $M - e$ is not connected (i.e. 2-connected), it has a 1-separation, i.e. there is $\emptyset \neq X \subsetneq E - e$ such that $\text{rk}(X) + \text{rk}(E - e - X) = \text{rk}(M - e) = \text{rk}(M)$, where the last equality holds since M is connected. Assuming without loss of generality that $|X| \leq |E - (X + e)|$, we have $\text{rk}(X + e) + \text{rk}(E - e - X) \leq$

$\text{rk}(M) + 1$, but since M is 3-connected, $X + e$ cannot be a 2-separation. Hence we must have $|X| = |E - (X + e)| = 1$ which implies $|E| = 3$. Now, the rank of $M - e$ cannot be 0 or 1 since $M - e$ is not connected but it does not have loops (or M would have loops). Hence $\text{rk}(M - e) = 2$ i.e. $M - e = U_{2,2}$. Then we have that, if $E = \{a, b, e\}$, $\text{rk}(M) = 2$ and $\{a, b\}$ is a basis. Moreover, $\{a, e\}$ without loss of generality is a basis (since e is not a loop) but then $\{b, e\}$ must be a basis as well, or a would be a coloop of M . \square

Lemma 4.36. *Let M be a connected matroid, F , C a flacet and circuit of M respectively, such that $|F| \geq 2$ and $F \cap C = \{e\}$ for some $e \in M$. Then $M - e$ is connected.*

Proof. Let M, F, C, e as in the hypotheses and assume by contradiction that $M - e$ is not connected. If M is 3-connected, $M = U_{3,2}$ thanks to Observation 4.35. But then we would get a contradiction as uniform matroids do not have flacets which are not singletons (see Observation 2.32). Hence M is not 3-connected and can be written as a 2-sum of 3-connected matroids: $M = M_1 \oplus_2 \cdots \oplus_2 M_k$, with $k > 1$. Among all counterexamples M to the theorem, take one with k minimum. Let M_1 be the unique matroid among M_1, \dots, M_k such that $e \in E(M_1)$. From Observation 4.34 and from the fact that if $M = M' \oplus_2 M''$, M is connected if and only if M', M'' are (Proposition 2.20), we derive that $M_1 - e$ is not connected, hence $M_1 = U_{3,2}$. This implies that M_1 has three elements, say a, b, e , not all of which are in $E(M)$ as they are deleted by 2-sums. Now, let T be the tree decomposition of M according to Theorem 2.33 and let v be the vertex of T corresponding to M_1 . v has degree at most 2. First suppose that v has degree 1, i.e. it is a leaf. Then we can write $M = U_{3,2} \oplus M'$ for some appropriate matroid M' , with $b \in E(M')$. But then either $a \notin C$, which implies by Observation 4.34 that $\{e, b\}$ is a circuit of M_1 (a contradiction), or $a \notin F$. In the latter case by Theorem 2.30 (and since $\{e\} \subsetneq F$) we must have that $\{e, b\}$ is a flacet of $U_{3,2}$, in contradiction with Observation ???. Now suppose that v has degree 2, i.e. there are M', M'' such that $M = M' \oplus_2 M''$, $E(M') \cap E(M'') = \{b\}$ (i.e. $b \notin E(M)$) and v is a leaf in the decomposition tree of M' . We apply Theorem 2.30 to M and consider several cases for F :

1. $F = E(M') \cup F'' - b$, where F'' is a flacet of M'' containing b . But then $C \subset E(M'') + e$ (since $C \cap F = \{e\}$ and $E(M') - b \subset F$), which implies by Observation 4.34 that $\{e, b\}$ is a circuit of M' , hence of M_1 , a contradiction.
2. $F = E(M'') \cup F' - b$, where F' is a flacet of M' containing b . In this case, $C \subset E(M') - b$ is a circuit of M' , $|F'| \geq 2$ since $e, b \in F'$, and $F' \cap C = \{e\}$. But then the hypotheses of the theorem are satisfied by M', F', C , contradicting to the minimality of k .
3. F is a flacet of M' not containing b . But then the hypotheses of the theorem are satisfied by M', F, C , again a contradiction.
4. $F = E(M') - b$: this implies as before that $C \cap E(M') + b = \{e, b\}$ is a circuit of M' , a contradiction.

□

Lemma 4.37. *Let $k \geq 3$ and $K = \{B_1, \dots, B_k\}$ a maximal clique in H , and let e_1, \dots, e_k the corresponding elements of M from Lemma 4.32. Let R_K be the set of rows r of S such that $S(r, B_i) = 0$ for exactly one $i \in [k]$, and R'_K be the set of rows r' of S such that $S(r', B_i) = 1$ for exactly one $i \in [k]$, and $r' \neq \mathbf{1} - r$ for all $r \in R_K$. Then the following holds:*

1. *The rows of $R_K \cup R'_K$ correspond to element inequalities.*
2. *$|R_K \cup R'_K| = k$ and in particular there is exactly an inequality for each e_i , i.e. each row of $R_K \cup R'_K$ corresponds to an inequality $x(e_i) \geq 0$ or $x(e_i) \leq 1$ for a different i .*
3. *Case 1 of Lemma 4.32 holds if and only if the rows of R_K are of the form $x(e) \geq 0$ and the rows of R'_K are of the form $x(e) \leq 1$. Similarly, case 2 of Lemma 4.32 holds if and only if the rows of R_K are of the form $x(e) \leq 1$ and the rows of R'_K are of the form $x(e) \geq 0$.*

Proof. 1. First, let $r \in R_K$ and assume by contradiction that r is not an element inequality, hence it is of the form $x(F) \leq \text{rk}(F)$ with $F \in \mathcal{F}$ and $|F| \geq 2$. At the cost of renaming B_1, \dots, B_k , we have $|F \cap B_1| = \text{rk}(F)$ and $|F \cap B_i| = \text{rk}(F) - 1$ for $i = 2, \dots, k$. We consider the two cases of Lemma 4.32:

- Case 1 holds, i.e. e_1, e_2, \dots, e_k form a circuit. But then for any $i = 2, \dots, k$, $e_i = B_1 \setminus B_i \in E(F)$ and by Observation 4.33, fact 1, this implies $e_1 \in E(F)$, which yields $|B_1 \cap E(F)| = |B_i \cap E(F)|$ for any $i = 2, \dots, k$, a contradiction.
- Case 2 holds, i.e. e_1, e_2, \dots, e_k form a cocircuit D . But then $e_1 = B_1 \setminus B_2 = \dots = B_1 \setminus B_k \in F$, and for $i = 2, \dots, k$ $e_i \notin F$, i.e. $D \cap F = \{e_1\}$, in contradiction with Observation 4.33, fact 3.

This proves that every $r \in R_K$ is an element row. Now assume that there is r' in R'_K of the form $x(F) \leq \text{rk}(F)$ with $F \in \mathcal{F}$ and $|F| \geq 2$. This means that without loss of generality $|B_1 \cap E(F)| = \text{rk}(F) - 1$, $|B_i \cap E(F)| = \text{rk}(F)$ for $i = 2, \dots, k$. We again distinguish two cases, following Lemma 4.32.

- Case 1 holds and $C = \{e_1, \dots, e_k\}$ is a circuit. But then $E(F) \cap C = \{e_1\}$, hence by Lemma 4.36 we have that e_1 is deletable, hence $\mathbf{1} - r \in R_K$, a contradiction.
- Case 2 holds, i.e. $D = \{e_1, e_2, \dots, e_k\}$ is a cocircuit. We then have that $D \cap F = D \setminus \{e_1\}$. Consider the dual matroid M^* , of which D is a circuit. Moreover, $E - F$ is a facet of M^* (Observation 4.33, fact 4), and $(E - F) \cap D = \{e_1\}$, which by Lemma 4.36 implies that $M^* - e_1$ is connected, but $M^* - e_1 = (M/e_1)^*$ and since a matroid is connected if and only if its dual is, we deduce that e_1 is contractible. But then $x(e_1) \leq 1$ is a row of R_K that is opposite to r , a contradiction.

2. This is an easy consequence of the previous statement, of the fact that every element in a connected matroid is either deletable or contractible (Lemma 4.27), and of the fact that for any $e \in E$, either $e = e_i$ for some i or e is in all B_i 's or in none of them.
3. First, suppose by contradiction that $r_1, r_2 \in R_K$ correspond to $x(e_1) \geq 0$, $x(e_2) \leq 1$ respectively. This means that, without loss of generality, $e_1 \in \cap_{i=2}^k B_i \setminus B_1$, and $e_2 \in B_2 \setminus \cup_{i \neq 2} B_i$. But this contradicts both case 1, 2 from Lemma 4.32. This argument shows the following: all the rows of R_K are of the same form, all the rows of R'_K are of the same form, and a row of R_K and a row of R'_K are of different forms. Now, the statement follows trivially.

□

We are now ready to outline an algorithm to produce $\mathcal{C}(M)$. As a first step, for any triangle Δ in H , we obtain the unique maximal clique K containing Δ and the corresponding R_K, R'_K . Note that we can assume that M contains at least a circuit of size at least 3, otherwise the fact that M is connected implies $M = U_{n,1}$: in the latter case $B(M)$ is affinely isomorphic to the $n - 1$ -dimensional simplex and its slack matrix is (a permutation of) the identity matrix, which can easily be recognized. Hence H contains at least a triangle, i.e. we obtain at least one maximal clique K , and moreover every element row of S is in some R_K or R'_K : indeed, it is easy to see that in a connected matroid that is not $U_{n,1}$ every element belongs to a circuit of size at least 3. Hence we obtain the element rows of S as a set $R_E = \{r \text{ row of } S : r \in R_K \cup R'_K \text{ for some maximal clique } K\}$. We now need to determine, for each of those rows, whether it has form $x(e) \geq 0$ or $x(e) \leq 1$. We will achieve this once we deal with the circuits of M that have size 2. We first observe the following:

Observation 4.38. *Let $C = \{e, f\}$ a circuit of M , M connected. Then there are two bases B_1, B_2 such that $B_1 + e = B_2 + f$ and for any facet F of M we have $|B_1 \cap F| = |B_2 \cap F|$. Moreover, the columns in S corresponding to B_1, B_2 differ exactly in the element rows relative to e and f .*

Proof. As M is connected, e is contained in a basis B_2 (which does not contain f) hence $B_1 = B_2 + f - e$ is a basis: indeed, if it is not there is a circuit $C' \subseteq B_1$ containing f , but then applying the circuit axiom $C \cup C' - f$ is a dependent set contained in B_2 , a contradiction. Now if there is a facet F such that $|B_1 \cap F| \neq |B_2 \cap F|$, this implies that $e \in F, f \notin F$ without loss of generality. But this is a contradiction to Observation 4.33, fact 1. The rest follows. □

Notice that the reverse of the previous statement is not true. So in the second step of our algorithm we go through all the edges of H and create a family \mathcal{C}_2 which contains subsets of rows of S which satisfy the conditions of the statement, hence capturing all the circuits of size 2 and possibly some other set. Formally, $\rho = \{r_1, \dots, r_k\}$, with $2 \leq k \leq 4$, is in \mathcal{C}_2 if and only if the following three conditions are satisfied: $\rho \subset R_E$; there is an edge B_1, B_2 of H such that the columns B_1, B_2 differ exactly in correspondence of $\{r_1, \dots, r_k\}$; and there are two rows in ρ , without loss of generality r_1, r_2 , so that any other row in ρ is opposite to r_1 or r_2 (for

instance if $k = 4$ we have that $r_3 = \mathbf{1} - r_1, r_4 = \mathbf{1} - r_2$). From Observation 4.38 we have that for any circuit $C = \{e, f\}$ of M , the element rows corresponding to e, f (which are at most four) are in a set of \mathcal{C}_2 . Moreover, for any $\rho \in \mathcal{C}_2$, we have that as soon as the form of one of the r_i is determined, the form of the other rows in ρ is determined as well: for instance, if $r_1, r_2 \in \rho$ are not opposite rows and $S(r_1, B_1) = 1, S(r_2, B_1) = 0$, we must have that r_1, r_2 have the same form, so if r_1 corresponds to $x(e) \geq 0$, r_2 must correspond to $x(f) \geq 0$. Notice that the same holds for R_K, R'_K , thanks to Lemma 4.37: once we determined whether K is inducing a circuit or a cocircuit, we know the form of each row in $R_K \cup R'_K$. We now use this in the final step of our algorithm.

Let \mathcal{C} be the family of subset of rows of S consisting of all the R_K 's, R'_K 's and of \mathcal{C}_2 , then every circuit of M corresponds to some set in \mathcal{C} . Our goal is to determine which subsets actually correspond to circuits (and which to cocircuits). Since as already noticed $B(M)$ and $B(M^*)$ have the same slack matrix, and the circuits of one are the cocircuits of the other, we can just fix any of the cliques and assume without loss of generality that it induces a circuit: this will determine the form of the rows in the corresponding R_K, R'_K , thanks to Lemma 4.37. Now, those rows will belong to other sets in \mathcal{C} , hence by determining their form we will determine the form of the other rows in those sets as well. We now argue that this propagates to all circuits of M . Let \mathcal{C}_M be the graph with $\mathcal{C}(M)$ as vertex set and where two circuits are adjacent if they share at least one element. Since M is connected, \mathcal{C}_M is connected. Therefore the choice on one circuit will eventually lead to identify the form of all the element rows. This gives us automatically all the circuits of M of size at least 3 (again by Lemma 4.37), and for the others we can just check for any $e, f \in E$ (corresponding to a subset in \mathcal{C}_2) whether $\{e, f\}$ is independent (i.e. whether there exists a basis $B \supset \{e, f\}$, which we can check once we know the element inequalities for e, f). This completes the construction of $\mathcal{C}(M)$. Below, we summarize the main steps of the algorithm.

Algorithm 1:

- 1 **for** Δ triangle of H **do**
 - 2 | Get K maximal clique containing K , and R_K, R'_K
 - 3 $R_E = \{r \text{ row of } S : r \in R_K \cup R'_K \text{ for some } K\}$
 - 4 Get \mathcal{C}_2 , and set $\mathcal{C} = \{R_K, R'_K \text{ for } K \text{ max clique}\} \cup \mathcal{C}_2$
 - 5 Choose any K and assign to rows in R_K the form $x(e) \geq 0$, and to rows in R'_K the form $x(e) \leq 1$
 - 6 Propagate the assignments until the form of all R_E has been determined
 - 7 **return** $\mathcal{C}(M)$
-

4.5.2 Phase 2: reconstructing $B(M)$ and its slack matrix

We will now show that we can efficiently recognize the slack matrix of $B(M)$ when M is assumed to be 2-level and connected, once we have the list L which, in the YES case, is equal to $\mathcal{C}(M)$. To do this we will essentially reconstruct M , its bases and its flacets, compute the slack matrix

of $B(M)$ and check whether it is equal to S . In this phase we will use the 2-levelness of $B(M)$, in particular the results we obtained in Section 2.4.

First, notice that from the previous discussion it follows that \mathcal{L} has size polynomial in n , the number of vertices of H (and of columns of S), as there are at most as many circuit as there are triangles in H . Now, assuming that \mathcal{L} is the list of the circuits of \mathcal{M} , this gives us an independence oracle for M , and with that we can enumerate all the bases of M in total polynomial time (folklore, see for instance [63]). This implies that we either find $n + 1$ bases and answer NO (since \mathcal{L} and S are not coherent) or we find exactly m bases, in both cases in polynomial time. Once we have the bases (i.e. the vertices of $B(M)$) we need to find the facet defining inequalities of $B(M)$ and check if the resulting slack matrix is equal to S , in which case we answer YES.

Since checking whether a matroid is connected can be done efficiently (using for instance the shifting algorithm given in [9]), we can check which element inequalities are facet defining (and whether this is coherent with the rows in R_E). We now argue that finding the flacets of our matroids can be done efficiently.

Recall from Theorem 2.33, 2.21 that, since $B(M)$ is 2-level and M is connected, M can be obtained from some uniform matroids U_1, \dots, U_k by a series of 2-sum operations, which are represented by a tree T . Also, Theorem 2.35 gives a simple description of the (linearly many) flacets of M in terms of cuts of T . To obtain T and U_1, \dots, U_k , one has to decompose M in smaller 3-connected matroids, again using the shifting algorithm. More precisely, we start from M , we run the algorithm and either obtain that M is 3-connected, or get M_1, M_2 such that $M = M_1 \oplus_2 M_2$, and repeat the algorithm on M_1, M_2 . This procedure can be done efficiently for any matroid, and moreover at the same time we can get a list of the bases of the smaller matroids: given $B \in \mathcal{B}(M)$ and M_1, M_2 on ground set E_1, E_2 , $B \cap E_1$ is a basis of M_1 if it has size $\text{rk}(M_1)$, and $B \cap E_1 + p$ is a basis otherwise, where p is the element so that $E_1 \cup E_2 = E + p$.

In this way, once obtained the 3-connected matroids whose 2-sum is M , and their tree structure, we can check that each of them is uniform by verifying that they have the right number of bases. We can hence verify that M is 2-level and then check, for any of the (linearly many) sets described in Theorem 2.35, whether they are flacets, and in this case whether a corresponding row is present in S . To conclude, we described a polynomial algorithm to recognize whether a given 0/1 matrix is the slack matrix of $B(M)$, for M 2-level and connected. Together with Theorem 4.22 and Lemma 4.28, which deal with the case where M is not connected, this gives an alternative proof of Theorem 4.31.

Remark 4.39. We conclude the section by remarking that this proof of Theorem 4.31 can be extended to recognize slack matrices of general matroid polytopes, on the assumption that the flacets of such matroids can be enumerated in total polynomial time. Indeed, the results in Section 4.5.1 hold for any matroid (notice that the element rows are 0/1 even in slack matrices of non-2-level matroid polytopes), and in Section 4.5.2 the only part where we explicitly use the 2-levelness of $B(M)$ is to efficiently obtain the flacets of M and reconstruct the slack matrix.

Chapter 4. Slack matrices of 2-level polytopes: recognition and decomposition

Moreover, as already remarked, while we define the 1-sum only on 0/1 matrices, it can be easily extended to general matrices, as well as all the theorems on the 1-sum. We are not aware of any result on enumeration of facets of the kind that is known for bases and circuits (see [63]).

5 Extended formulations in output-efficient time from communication protocols

5.1 Introduction

Yannakakis' Theorem (Theorem 1.5, see also [100]) implies that an extended formulation for a polytope can be constructed via a nonnegative factorization of its slack matrix. While constructive, this result is not output-efficient, since the time needed to produce the extended formulation does not only depend on the size of the formulation itself but also on the size of the original description of the polytope. In this chapter we deal with this problem in particular with respect to factorizations obtained via communication protocols. Most notably, we give sufficient conditions under which a deterministic communication protocol can be turned in an algorithm to write an extended formulation in time linear in the size of such formulation. The most famous example of the application of deterministic protocols to extended formulations is Yannakakis' protocol [100], that implies the existence of an extended formulation of quasipolynomial size for the stable set polytope of perfect graphs. Our original motivation for this work was to make Yannakakis' result output-efficient: we achieve this by giving two different algorithms that produce the desired extended formulation in quasipolynomial time, one as a consequence of a more general theorem and one as a direct and more efficient construction. We also obtain compact formulations for the stable set polytope of some subclasses of perfect graphs.

Contribution and organization. The chapter is organized as follows:

- In Section 5.2 we describe the connection between communication protocols and extension complexity from [100], the concept of extended formulation of a pair of polytopes from [81], and some notions on the stable set polytope and its clique relaxation. We also clarify the scope of applicability of our results and the assumptions we make in order to make them efficient, in particular regarding Theorem 5.6.
- In Section 5.3 we give a simple general procedure to construct extended formulations from deterministic protocols, see Theorem 5.6. We then show the applicability of our result, by deriving an extended formulation for the stable set polytope of perfect graphs.

- In Section 5.4, we show that, in interesting special cases, one can derive an explicit, compact formulation by ad-hoc arguments, without relying on Theorem 5.6. In particular we give an alternative formulation for $STAB(G)$, G perfect, that has significant advantages over the previous one in terms of efficiency and applicability. We also give formulations the stable set polytope of claw-free perfect graphs and of comparability graphs.

5.2 Preliminaries

5.2.1 Deterministic and non-deterministic protocols

We start by describing the general setting of communication complexity. For a more detailed description we refer to [70]. Let M be a matrix with row (resp. column) set X (resp. Y). Consider two agents, Alice and Bob, who aim at computing the matrix M under partial information. In particular Alice is given as input a row index $i \in X$, Bob a column index $j \in Y$, and they aim at determining M_{ij} by exchanging information according to some pre-specified mechanism, that goes under the name of *protocol*. The protocol that they follow is said to *compute* M if, for any input i of Alice and j of Bob, it returns M_{ij} ; it is *deterministic* if the actions of Alice (resp. Bob) at any given step only depend on her (resp. his) input and on what they exchanged so far. The *complexity* of such a protocol is the maximum amount of bits exchanged in any execution. Such a protocol can be modelled as a rooted tree, with each vertex modelling a step where one of Alice or Bob speaks (hence labelled with A or B), and its children representing subsequent steps according to the different messages that can be sent at that stage. The *leaves* of the tree indicate the termination of the protocol and are labelled with the corresponding output. Assuming, without loss of generality, that each message exchanged consists of a single bit, we obtain that the tree is binary, with each edge representing a 0 or a 1 sent. Hence, a deterministic protocol can be identified by the following parameters: a rooted binary tree τ with node set \mathcal{V} ; a function $\ell : \mathcal{V} \rightarrow \{A, B\}$ (“Alice”, “Bob”) associating each vertex to its type; for each leaf $v \in \mathcal{V}$, a positive number λ_v corresponding to the value output at v ; for each $v \in \mathcal{V}$ that is not a leaf and such that $\ell(v) = A$ (resp. $\ell(v) = B$) the set of inputs $S_v \subseteq X$ (resp. $S_v \subseteq Y$) such that Alice (resp. Bob) sends a 1 at node v . We represent this succinctly by $(\tau, \ell, \lambda, \{S_v\}_{v \in \mathcal{V}})$.

An execution of the protocol corresponds to a path of τ from the root to a leaf, whose edges correspond to the bits sent during the execution. The set of input indices (i, j) that produce the same execution, i.e. leading to the same leaf v , correspond to entries of M with the same value λ_v , and moreover, without loss of generality, they can be assumed to form a submatrix of M : indeed, at the end of the protocol, both Alice and Bob can be assumed to know the outcome, each independently of the input of the other (see [70] for more details). Such submatrices with constant value are called *monochromatic rectangles*.

The complexity of the protocol is given by the height h of the tree τ . Hence a deterministic protocol computing M gives a partition of M in at most 2^h monochromatic rectangles. We

remark that one can obtain a protocol (and a partition in rectangles) for M^T given a protocol for M by just exchanging the roles of Alice and Bob.

The setting of *non-deterministic* protocols is similar as before, but now Alice and Bob are allowed to make guesses in their communication, with the requirement that, at the end of the protocol, they can both independently verify that the outcome corresponds to M_{ij} for at least one guess made during the protocol. A nondeterministic protocol is called *unambiguous* if for any input i, j , exactly one guess allows to verify the value of M_{ij} . The *complexity* of a nondeterministic protocol is the maximum (over all inputs and guesses) amount of bits exchanged during the protocol. Nondeterministic protocols of complexity c provide a cover of M with at most 2^c monochromatic rectangles, which is a partition in the case the protocol is unambiguous. Moreover, each partition of M in N monochromatic rectangles corresponds to an unambiguous protocol of complexity $\lceil \log_2 N \rceil$, where Alice guesses the rectangle covering i, j .

We want to mention another class of communication protocols that is relevant to extended formulations, namely *randomized* protocols that compute a (nonnegative) matrix in expectations. These generalize both deterministic and nondeterministic protocols and have been defined in [30], where they are shown to be equivalent to non-negative factorizations (see the next section) and to essentially capture the notion of extension complexity. However, our results do not extend to such general protocols, hence we do not formally define them and we refer the interested reader to [30].

5.2.2 Extended formulations for a pair of polytopes

Yannakakis' Theorem has been extended multiple times and generalized (see [30, 35, 81]). In particular, in [81] the concept of extended formulation is applied to a pair (P, Q) , where P, Q are polytopes with $P \subseteq Q \subseteq \mathbb{R}^d$ where P is given as the convex hull of vertices and Q via a set of linear inequalities. A polyhedron $Q \in \mathbb{R}^{d'}$ is an extension for the pair (P, Q) if there is a projection $\pi : \mathbb{R}^{d'} \rightarrow \mathbb{R}^d$ such that $P \subseteq \pi(Q) \subseteq Q$. The concepts of extended formulation and extension complexity of a pair are defined analogously as in Chapter 1, and clearly these new definitions coincide with the previous ones if $P = Q$.

Definition 5.1. Given a polytope $Q = \text{conv}(v_1, \dots, v_n) \subseteq \mathbb{R}^d$ and a polyhedron $Q = \{x \in \mathbb{R}^d : Ax \leq b\}$, where A has m rows, the *slack matrix* $M(P, Q)$ of the pair (P, Q) is a non-negative $m \times n$ matrix with $M(P, Q)_{i,j} = b_i - a_i^\top v_j$, i.e., the (i, j) -th entry is the slack of point v_j of P with respect to the i -th inequality in the description of Q .

Given a non-negative matrix $M \in \mathbb{R}_{\geq 0}^{m \times n}$, a *non-negative factorization* of M is an expression of the form $M = TU$, where T, U are non-negative matrices. Recall from Chapter 1 that the non-negative rank of M is the smallest intermediate dimension in a non-negative factorization of M , and that the extension complexity of a polytope is equal to the non-negative rank of its slack matrix (Theorem 1.5). This has been generalized to pairs of polytopes in [81]. We report

Chapter 5. Extended formulations in output-efficient time from communication protocols

below a version of the latter, adapted to our purposes.

Theorem 5.2 (Yannakakis' Theorem for pairs of polytopes). *Given a slack matrix M of a pair of polytopes (P, Q) of dimension at least 1, the extension complexity of (P, Q) is equal to the nonnegative rank of M . In particular, if $M = TU$ is a non-negative factorization of M , then*

$$P \subseteq \{x : \exists y \geq 0 : Ax + Ty = b\} \subseteq Q. \quad (5.1)$$

Hence, a factorization of the slack matrix of intermediate dimension N gives an extended formulation of size N (i.e. with N inequalities). However such formulation has as many equations as the number of rows of A . While at most N of these equation are non-redundant, there is no clear a priori way of reducing the system of equations without listing all of them.

5.2.3 Protocols and extended formulations

Assume we have a deterministic or an unambiguous protocol of complexity c for computing a slack matrix M of a polytope P (or equivalently of a pair (P, Q)). We assume for simplicity that M is a 0/1 matrix, but our arguments extend to the general case. As described above, the protocol gives a partition of M into at most 2^c monochromatic rectangles. This implies that $M = R_1 + \dots + R_N$, where $N \leq 2^c$ and each R_i is a rank 1 matrix corresponding to a 1-rectangle (a monochromatic rectangle of value 1). Hence M can be written as a product of two non-negative (0/1) matrices U, T of intermediate dimension N , where $T_{i,j} = 1$ if the rectangle R_j contains row index i , and $U_{i,j} = 1$ if R_j contains column index j . As a consequence of Theorem 5.2, this yields an extended formulation for P . In particular, let $P = \{x \in \mathbb{R}^d : Ax \leq b\}$, with $A \in \mathbb{R}^{m \times d}$, let \mathcal{R} be the set of 1-rectangles of M , and, for $i = 1, \dots, m$, let a_i be the i -th row of A and $\mathcal{R}_i \subset \mathcal{R}$ be the set of rectangles whose row index set includes i . Then the following is an extended formulation for P :

$$\begin{aligned} a_i x + \sum_{R \in \mathcal{R}_i} y_R &= b_i \quad \forall i = 1, \dots, m \\ y &\geq 0 \end{aligned} \quad (5.2)$$

Again, the formulation has as many equations as the number of rows of A , and it is not clear how get rid of non-redundant equations. Here it is important to address the issue of what is our input, and what assumptions we need in order to get an "efficient" algorithm. The following discussion is not formal and has the purpose to explain the applicability of the results of this chapter, while the formal details will be clarified in the next section.

Recall that, in our setting, the matrix A describing P is thought as being exponential in size, while $|\mathcal{R}|$ is polynomial (or quasipolynomial). We assume that we have an *implicit* representation of our polytope P of interest, and in particular of A . This assumption is natural as, without it, we can hardly imagine to have any useful protocol for the slack matrix of P . As

an example, consider the case, discussed below, of the stable set polytope of perfect graphs, of which we know the vertices and inequalities without of course having to explicitly list them (as that would take exponential time).

Recall that a deterministic protocol is identified with a tuple $\tau, \ell, \Lambda, \{S_v\}_{v \in \mathcal{V}}$. While we can assume that τ, ℓ, Λ are given to us explicitly, the sets S_v have in general exponential size. Hence we assume to have an implicit description of them, in particular of our rectangles \mathcal{R} : notice that the latter correspond to leaves of τ and can be identified by a sequence of bits exchanged during the protocol. Knowing the structure of our protocol gives us an implicit representation of \mathcal{R}_i for each i : again, this is a reasonable and basic assumption for approaching the formulation (5.2) from an algorithmic point of view.

Now, the natural approach to reduce the size of (5.2) is to eliminate redundant equations. However the structure of the coefficient matrix depends both on A and on the sets \mathcal{R}_i 's, which can have a very complex behaviour. To get a better understanding of the issue the reader is encouraged to try on the example of $\text{STAB}(G)$, G perfect: the sets \mathcal{R}_i 's have very non-trivial relations with each other that depend heavily on the graph, and (although one can exploit some structure as we will see at the end of Section 5.3) we did not manage to directly reduce the system (5.2) for general perfect graphs. Theorem 5.6 shows how to bypass this problem for any deterministic protocol. Informally, we shift the problem of eliminating redundant equations from the system (5.2) to a family of systems $\{A_R x + y_R = b_R, y_R \geq 0\}$, one for each rectangle R produced by the protocol, where y_R is a single variable. The latter systems can still have exponential size, but they are in principle much easier to deal with since their structure only depends on (a submatrix of) A .

5.2.4 The stable set polytope

The most famous application of protocols to extended formulations is probably the first one, proved in [100], in the context of stable set polytopes. We recall that for general graphs, $\text{STAB}(G)$ has exponential extension complexity [35, 42] and no "explicit" linear description of it is known. The *clique relaxation* of $\text{STAB}(G)$ is the following:

$$\text{QSTAB}(G) = \left\{ x \in \mathbb{R}_+^d : \sum_{v \in C} x_v \leq 1 \text{ for all cliques } C \text{ of } G \right\},$$

Notice that, in the above description, one can restrict to maximal cliques, even though in the following we will consider all cliques whenever it is convenient. As a consequence of the equivalence between separation and optimization [50], optimizing over $\text{QSTAB}(G)$ is NP-hard for general graphs. However, the clique relaxation is exact for perfect graphs, for which the optimization problem is polynomial time solvable (see Chapter 1):

Theorem 5.3 ([16]). *A graph G is perfect if and only if $\text{STAB}(G) = \text{QSTAB}(G)$.*

The following result, from [100], has been mentioned in Chapters 1 and 3. Here we state it in a

Chapter 5. Extended formulations in output-efficient time from communication protocols

more convenient form.

Theorem 5.4. *Let G be a graph with n vertices. There is a deterministic protocol of size $O(\log^2 n)$ computing the slack matrix of the pair $(STAB(G), QSTAB(G))$. In particular, there is an extended formulation of size $n^{O(\log(n))}$ for $(STAB(G), QSTAB(G))$.*

We remark that, when G is perfect, Theorem 5.4 gives a quasipolynomial size extended formulation for $STAB(G)$. At the end of Section 5.3 we give a modified version of this protocol.

5.3 A general approach

Let us start by recalling the well-known theorem from Balas [4], in a version given by Weltge ([98], Section 3.1.1).

Theorem 5.5. *Let $P_1, P_2 \subset \mathbb{R}^d$ be polytopes, with $P_i = \pi_i\{y \in \mathbb{R}^{m_i} : A^i y \leq b^i\}$, where $\pi_i : \mathbb{R}^{m_i} \rightarrow d$ is a linear map, for $i = 1, 2$. Let $P = \text{conv}(P_1 \cup P_2)$. Then we have:*

$$P = \{x \in \mathbb{R}^d : \exists y^1 \in \mathbb{R}^{m_1}, y^2 \in \mathbb{R}^{m_2}, \lambda \in \mathbb{R} : x = \pi_1(y^1) + \pi_2(y^2), \\ A^1 y^1 \leq \lambda b^1, A^1 y^2 \leq (1 - \lambda)b^2, 0 \leq \lambda \leq 1\}.$$

Moreover, the inequality $\lambda \geq 0$ ($\lambda \leq 1$ respectively) is redundant if P_1 (P_2) has dimension at least 1. Hence

$$xc(P) \leq xc(P_1) + xc(P_2) + |\{i : \dim(P_i) = 0\}|.$$

We now give our general theorem to efficiently turn deterministic protocols into explicit extended formulations. Its proof is inspired by [32], where a general method is given to construct extended formulations for polytopes specified by boolean formulas. While similar in flavour, it seems that these two results are incomparable, in the sense that one does not follow from the other. It is possible however that they both fall under a more general framework which has not been investigated yet.

Note that the following result relies on the existence of a deterministic protocol $(\tau, \ell, \Lambda, \{S_v\}_{v \in \mathcal{V}})$, but its complexity does not depend on the encoding of Λ and $\{S_v\}_{v \in \mathcal{V}}$ (see the previous section).

Theorem 5.6. *Let S be a slack matrix for a couple (P, Q) , where $P = \text{conv}\{x_1^*, \dots, x_n^*\} \subseteq \mathbb{R}^d$ and $Q = \{x \in \mathbb{R}^d : a_i x \leq b_i \text{ for } i = 1, \dots, m\}$ are polytopes and for $j \in [d]$, let ℓ_j (resp. u_j) be a valid upper bound (resp. lower bound) on variable x_j in Q . Assume there exists a deterministic protocol $(\tau, \ell, \Lambda, \{S_v\}_{v \in \mathcal{V}})$ with rectangle set \mathcal{R} and complexity $c \leq \lceil \log_2 |\mathcal{R}| \rceil$ computing S . For any $R \in \mathcal{R}$, let $P_R = \text{conv}\{x_j^* : j \text{ is a column of } R\}$ and $Q_R = \{x \in \mathbb{R}^d : a_i x \leq b_i \forall i \text{ row of } R; \ell_j \leq x_j \leq u_j \text{ for all } j \in [d]\}$.*

Suppose we are given τ, ℓ and for each $R \in \mathcal{R}$ an extended formulation T_R for (P_R, Q_R) . Let $\sigma(T_R)$ be the size (number of inequalities) of T_R , and $\sigma_+(T_R)$ be the total encoding length of the

description of T_R (including the number of inequalities, variables and equations). Then we can construct an extended formulation for (P, Q) of size linear in $\sum_{R \in \mathcal{R}} \sigma(T_R)$ in time linear in $\sum_{R \in \mathcal{R}} \sigma_+(T_R)$.

Proof. We can assume without loss of generality that τ is a complete binary tree, i.e. each node of the protocol other than the leaves has exactly two children. Let \mathcal{V} be the set of nodes of τ and $\nu \in \mathcal{V}$. Note that there exists exactly one (non-necessarily monochromatic) rectangle S_ν associated to ν , which is given by the pairs (i, j) such that, on input (i, j) , the execution of the protocol visits node ν . Let us define, for any such S_ν , a pair (P_ν, Q_ν) as follows:

$$P_\nu = \text{conv}\{x_j^* : j \text{ is a column of } S_\nu\} \text{ and}$$

$$Q_\nu = \{x \in \mathbb{R}^d : a_i x \leq b_i \ \forall \ i \text{ row of } S_\nu; \ \ell_j \leq x_j \leq u_j \text{ for all } j \in [d]\}.$$

Clearly, for any ν one has $P_\nu \subseteq P \subseteq Q \subseteq Q_\nu$. Notice that, if $\rho \in \mathcal{V}$ denotes the root of the protocol, we have $S_\rho = S$, $P_\rho = P$, and $Q_\rho = Q$. We now show how to obtain an extended formulation for the pair (P_ν, Q_ν) given extended formulations T_{ν_i} 's for (P_{ν_i}, Q_{ν_i}) , $i = 0, 1$, where ν_0 (resp. ν_1) are the two children nodes of ν .

Assume first that ν is labelled A . Then we have $S_\nu = \begin{bmatrix} S_{\nu_0} \\ S_{\nu_1} \end{bmatrix}$ (up to permutation of rows), and therefore $P_\nu = P_{\nu_0} = P_{\nu_1}$ and $Q_\nu = Q_{\nu_0} \cap Q_{\nu_1}$. Hence we have $P_\nu \subseteq \pi_0(T_{\nu_0}) \cap \pi_1(T_{\nu_1}) \subseteq Q_\nu$, where π_i is a projection from the space of T_{ν_i} to \mathbb{R}^d . An extended formulation for $T_\nu := \pi_0(T_{\nu_0}) \cap \pi_1(T_{\nu_1})$ can be obtained efficiently by juxtaposing the formulations of T_{ν_0}, T_{ν_1} .

Now assume that ν is labelled B . Then we have $S_\nu = \begin{bmatrix} S_{\nu_0} & S_{\nu_1} \end{bmatrix}$ (up to permutations of columns). Hence, $P_\nu = \text{conv}\{P_{\nu_0} \cup P_{\nu_1}\}$ and $Q_\nu = Q_{\nu_0} = Q_{\nu_1}$, which implies $P_\nu \subseteq \text{conv}\{\pi_0(T_{\nu_0}) \cup \pi_1(T_{\nu_1})\} \subseteq Q_\nu$. An extended formulation for $T_\nu := \text{conv}\{\pi_0(T_{\nu_0}) \cup \pi_1(T_{\nu_1})\}$ can be obtained efficiently by applying Theorem 5.5 to the formulations of T_{ν_0}, T_{ν_1} . Iterating this procedure, in a bottom-up approach we can obtain an extended formulation for (P, Q) from extended formulations of (P_ν, Q_ν) , for each leaf ν of the protocol.

We now bound the number of basic operations necessary to obtain our formulation. If $T_\nu = \pi_0(T_{\nu_0}) \cap \pi_1(T_{\nu_1})$, then $\sigma_+(T_\nu) \leq \sigma_+(T_{\nu_0}) + \sigma_+(T_{\nu_1})$. Consider now $T_\nu = \text{conv}\{\pi_0(T_{\nu_0}) \cup \pi_1(T_{\nu_1})\}$. From Theorem 5.5 we have $\sigma_+(T_\nu) \leq \sigma_+(T_{\nu_0}) + \sigma_+(T_{\nu_1}) + O(d)$. Now, since the binary tree associated to the protocol is complete, it has size linear in the number of leaves (it has actually $2|\mathcal{R}| - 2$ vertices), hence the final formulation T_ρ satisfies

$$\sigma_+(T_\rho) \leq \sum_{R \in \mathcal{R}} \sigma_+(T_R) + O(d) = O\left(\sum_{R \in \mathcal{R}} \sigma_+(T_R)\right),$$

where the last equation is justified by the fact that we can assume $\sigma_+(T_R) \geq d$ for any R . The bound on the size of T_ρ is derived in an analogous way. \square

Chapter 5. Extended formulations in output-efficient time from communication protocols

We remark that the theorem above does not guarantee that the formulation has *exactly* the form given by the one of the corresponding protocol. Also, let us note that, even for the special case $P = Q$, in the proof of the previous theorem we need the generalized version of Yannakakis' theorem for pairs of polytopes.

Last, observe that the proof of the previous theorem does not strictly require that we know extended formulations for nodes of the protocols corresponding to leaves. A similar bottom-up approach would work starting at any node v of the protocol, as long as we have an extended formulation for (P_v, Q_v) .

5.3.1 Application to $(STAB(G), QSTAB(G))$

We now describe how to apply Theorem 5.6 to the protocol from Theorem 5.4 as to obtain an extended formulation for $(STAB(G), QSTAB(G))$ in time $n^{O(\log(n))}$. In particular, this gives an extended formulation for $STAB(G)$, G perfect within the same time bound.

We first give a modified version of the protocol, stressing a few details that will be important in later sections. The reader familiar with the original protocol can immediately verify its correctness. Let v_1, \dots, v_n be the vertices of G in any order. At the beginning of the protocol, Alice is given a clique C of G as input and Bob a stable set S , and they want to compute the entry of the slack matrix of $STAB(G)$ corresponding to C, S , i.e. to establish whether C, S intersect or not.

At each stage of the protocol, the vertices of the current graph $G = (V, E)$ are partitioned between *low degree* L (i.e. at most $|V|/2$) and *high degree* H . Suppose first $|L| \geq n/2$. Alice sends (i) the index of the low degree vertex of smallest index in C , or (ii) 0 if no such vertex exists. In case (i), if $v_i \in S$, then $C \cap S \neq \emptyset$ and the protocol ends; else, G is replaced by $G \cap N(v_i) \setminus \{v_j \in L : j < i\}$, where $G \cap U$ denotes the subgraph of G induced by U . In case (ii), if Bob has no high degree vertex, then $C \cap S = \emptyset$ and the protocol ends, else, G is replaced by $G \cap H$. If conversely $|L| < n/2$, then the protocol proceeds symmetrically to above: Bob sends (i) the index of the high degree vertex of smallest index in S , or (ii) 0 if no such vertex exists. In case (i), if $v_i \in C$, then $C \cap S \neq \emptyset$ and the protocol ends; else, G is replaced by $G \cap \bar{N}(v_i) \setminus \{v_j \in H : j < i\}$. In case (ii), if Alice has no low degree vertex, then $C \cap S = \emptyset$ and the protocol ends, else, G is replaced by $G \cap L$. Note that at each step the number of vertices of the graph is decreased by at least half, and C and S do not intersect in any of the vertices that have been removed.

Now let S be the slack matrix of the pair $(STAB(G), QSTAB(G))$. Each rectangle R in which the protocol from Theorem 5.4 partitions S is univocally identified by the list of cliques and of stable sets corresponding to its rows and columns. With a slight abuse of notation, for a clique C (resp. stable set S) whose corresponding row is in R , we write $C \in R$ (resp. $S \in R$), and we also write $(C, S) \in R$. We let P_R be the convex hull of stable sets $S \in R$ and Q_R the set of clique inequalities corresponding to cliques $C \in R$, together with the unit cube constraints.

We need a fact on the structure of rectangles, for which we introduce some more notation: for a rectangle R , let C_R be the set of vertices sent by Alice and S_R the set of vertices sent by Bob during the corresponding execution of the protocol. Note that C_R is a clique and S_R is a stable set.

Observation 5.7. *There is exactly one clique C and one stable set S of G such that $C = C_R$ and $S = S_R$. Conversely, given a clique C and a stable set S , there is at most one rectangle R such that $C = C_R$ and $S = S_R$. Notice that $|C_R| + |S_R| \leq \lceil \log_2 n \rceil$ for any $R \in \mathcal{R}$.*

Now, assuming we are given the graph G as input, in order to apply Theorem 5.6 we need to perform two steps:

1. Obtain the tree T with label set ℓ deriving from the protocol for G .

A simple way is to first enumerate all cliques and stable sets of G of size at most $\lceil \log_2 n \rceil$ and run the protocol on each possible input pair to get \mathcal{R} (thanks to Observation 5.7). Then, derive the structure of T (and ℓ) from the rectangles obtained: for instance, all the rectangles whose C_R begins with vertex v_1 are descendants of the child of the root whose edge is labelled v_1 , etc.

2. For each leaf of T corresponding to a rectangle R , give a compact extended formulation T_R for the pair (P_R, Q_R) .

Fix $R \in \mathcal{R}$ to be a 1-rectangle, the 0-rectangle case being similar. Since R is a non-negative rank-1 matrix, an extended formulation of (P_R, Q_R) is given by

$$\{x \in \mathbb{R}^d, y_R \in \mathbb{R} : x(C) + y_R = 1 \ \forall C \in R, \ y_R \geq 0, 0 \leq x \leq 1\}. \quad (5.3)$$

We now reduce the equations in the description above, which can be exponentially many, to a smaller system. For that we need the following fact on the structure of the rectangles.

Lemma 5.8. *Let $R = (C_R, S_R)$ and $(C, S) \in R$. Then for any C' such that $C_R \subseteq C' \subseteq C$ and any S' such that $S_R \subseteq S' \subseteq S$ we have $(C', S') \in R$.*

Proof. Note that a vertex $v \in C \setminus C_R$ is not sent during the protocol on input (C, S) . Hence, the execution of the protocols on inputs (C, S) and $(C - v, S)$ coincides. Indeed at every step Alice chooses the first vertex of low degree in her current clique, and if v is never chosen, having v in the clique does not affect her choice. Moreover, the choice of Bob only depends on his current stable set and the vertices previously sent by Alice. In particular, we have $(C \setminus \{v\}, S) \in R$. Iterating the argument (and applying the symmetric for $v \in S \setminus S_R$) we conclude the proof. \square

Chapter 5. Extended formulations in output-efficient time from communication protocols

Now, we claim that

$$\begin{aligned}
 T_R = \{x \in \mathbb{R}^d, y_R \in \mathbb{R} : & x(C_R) + y_R = 1 \\
 & x(C_R + v) + y_R = 1 \quad \forall v \in V \setminus C_R : C_R + v \in R \\
 & y_R \geq 0 \\
 & 0 \leq x \leq 1\}.
 \end{aligned}$$

is an extended formulation for (P_R, Q_R) . It suffices to show that the coefficient vector of each equation from (5.3) is spanned by the coefficient vectors from equations in the formulation T_R above. Let $C \in R$. For any $v \in C \setminus C_R$, we have $C_R + v \in R$ thanks to Lemma 5.8. Hence we obtain:

$$\sum_{v \in C \setminus C_R} (x(C_R + v) + y_R) - (|C \setminus C_R| - 1)(x(C_R) + y_R) = x(C) + y_R,$$

as required.

We conclude by observing that the approach described above proceeds by obtaining the leaves of T , with an expensive enumeration of cliques and stable sets, and then it reconstructs T . This takes time $\Theta(n^{\lceil \log_2 n \rceil})$. However, one could instead try to construct T from the root, by distinguishing cases for each possible bit sent by Alice or Bob. This intuition is the basis for the alternative formulation that we give in the next section.

5.4 Direct derivations

5.4.1 Complement graphs

In order to derive an alternative formulation for $\text{STAB}(G)$, G perfect, we exploit the relationship between a perfect graph and its complement with respect to the stable set polytope. In this section, we show that an explicit formulation for $\text{STAB}(G)$, G perfect, can be easily obtained from an extended formulation of $\text{STAB}(\bar{G})$, keeping a similar size (including the number of equations).

The next Lemma can be found in [89, Section 65.4].

Lemma 5.9. *G is a perfect graph if and only if $\text{STAB}(G) = \{x : x \geq 0, x^T y \leq 1 \ \forall y \in \text{STAB}(\bar{G})\}$.*

The next Lemma is a restatement of Lemma 3.3 in the form of [98].

Lemma 5.10. *Given a non-empty polyhedron Q and $\gamma \in \mathbb{R}$, let $P = \{x : x^T y \leq \gamma \ \forall y \in Q\}$. If $Q = \{y : \exists z : Ay + Bz \leq b, Cy + Dz = d\}$, then we have that*

$$\begin{aligned}
 P = \{x : \exists \lambda \geq 0, \mu : & A^T \lambda + C^T \mu = x, \\
 & B^T \lambda + D^T \mu = 0, b^T \lambda + d^T \mu \leq \gamma\}.
 \end{aligned}$$

Hence $xc(P) \leq xc(Q) + 1$.

Now a straightforward calculation shows that, for a perfect graph G , $STAB(\bar{G})$ admits an extended formulation with approximately the same encoding length of an extended formulation of $STAB(G)$.

Corollary 5.11. *Let G be a perfect graph on n vertices such that $STAB(\bar{G})$ admits an extended formulation Q with r additional variables (i.e. $n + r$ variables in total), m inequalities and k equations. Then $STAB(G)$ admits an extended formulation with $m + k$ additional variables, $n + r + 1$ inequalities, $n + r$ equations, which can be written down explicitly given Q . In particular the size of such formulation is at most one plus the size of Q .*

Proof. This follows trivially from Lemma 5.10 and Lemma 5.9. The last sentence is justified by the fact that, since Q has at least one vertex, $m \geq n + r$. \square

5.4.2 Alternative formulation for $STAB(G)$, G perfect

We now present an algorithm that, given a perfect graph G on n vertices, produces an explicit extended formulation for $STAB(G)$ of size $n^{O(\log n)}$, in time bounded by $n^{O(\log n)}$. The algorithm is inspired by Yannakakis' protocol, even though the formulation obtained is different from what one would get from the factorization given by such protocol: the additional variables do not necessarily correspond to rectangles of the slack matrix.

Consider the protocol as described at the end of Section 5.3. Our algorithm can be seen as performing breadth-first search on the tree corresponding to the protocol for G , and iteratively decomposing G according to the non-leaf vertices met. When we meet a node v in which Alice speaks, we consider a different subgraph for each possible message sent (i.e. for each children of the node), and, as we will show, this corresponds to a partition of the clique constraints of the formulation of $STAB(G_v)$. When we meet a node in which Bob speaks, we would like to keep partitioning our constraints (even though Bob sends information about vertices of $STAB(G)$); hence we consider the complement of the current graph, in this way swapping cliques and stable sets, hence constraints and vertices, and the role of Alice and Bob, and proceed similarly as before. Notice that, in practice, we can stop exploring a branch as soon as we meet a subgraph that is small enough (or is a clique, an empty graph, or any graph for which we can efficiently get an extended formulation). When our search ends, we will go bottom-up by iteratively adding together the formulations obtained for the children and get a formulation for the parent (see Lemma 5.12), while using the construction given in Lemma 5.9 whenever a complement graph was taken, until we reach the root and obtain a formulation for $STAB(G)$.

The details and the proof of correctness of the algorithm are given below. We recall that, for a vertex v of G , $N^+(v) = N(v) + v$ denotes the inclusive neighbourhood of v .

Chapter 5. Extended formulations in output-efficient time from communication protocols

Lemma 5.12. *Let G be a perfect graph on vertex set $V = \{v_1, \dots, v_n\}$, and, fix k with $1 \leq k \leq n$. Let G_i be the induced subgraph of G on vertex set $V_i = N^+(v_i) \setminus \{v_1, \dots, v_{i-1}\}$ for $i = 1, \dots, k$, and G_0 the induced subgraph of G on vertex set $V_0 = \{v_{k+1}, \dots, v_n\}$. For $i = 0, \dots, k$ let $P_i = \text{STAB}(G_i) \times \mathbb{R}^{V \setminus V_i}$. Then we have*

$$\text{STAB}(G) = P_0 \cap \dots \cap P_k.$$

Proof. We first remark that, since by definition $v_i \in V_i$ for $i = 1, \dots, k$, and $V_0 = \{v_{k+1}, \dots, v_n\}$, we have $V_0 \cup \dots \cup V_k = V$. Recall that $\text{STAB}(G) = \{x \in \mathbb{R}^n : x(C) \leq 1 \text{ for } C \in \mathcal{C}, x \geq 0\}$, where \mathcal{C} is the set of cliques of G . Let us consider the partition $\mathcal{C}_0, \dots, \mathcal{C}_k$ of \mathcal{C} where a clique C is in \mathcal{C}_i , for $i \geq 1$, if $\{v_1, \dots, v_i\} \cap C = \{v_i\}$, and is in \mathcal{C}_0 otherwise. Let $x \in \mathbb{R}^n$. We have that $x \in P_0 \cap \dots \cap P_k$ if and only if $x_{V_i} \in \text{STAB}(G_i)$ for $i = 1, \dots, k$, with x_{V_i} denoting the restriction of x to coordinates in V_i . We now claim that $\text{STAB}(G_i) = \{x \in \mathbb{R}^{V_i} : x(C) \leq 1 \text{ for } C \in \mathcal{C}_i, x \geq 0\}$ for $i = 0, \dots, k$. Notice that $C \in \mathcal{C}_i$ implies that C is a clique in G_i , which proves the “ \subseteq ” inclusion. For the opposite inclusion, since each G_i , being an induced subgraph of G , is perfect, it suffices to show that \mathcal{C}_i includes all the maximal cliques of G_i . Let C be a maximal clique of G_i . If $i = 0$, then $C \in \mathcal{C}_0$ and we are done. If $i \geq 1$, then $v_i \in C$ as $V_i \subseteq N^+(v_i)$, and $C \in \mathcal{C}_i$, which concludes the proof of the claim. Hence we have that $x \in P_0 \cap \dots \cap P_k$ if and only if $x \geq 0$, and $x(C) \leq 1$ for any clique in $\mathcal{C}_0 \cup \dots \cup \mathcal{C}_k = \mathcal{C}$, but this is equivalent to $x \in \text{STAB}(G)$. \square

We now make a simple observation which is the basis of our decomposition approach.

Observation 5.13. *Let $P_1, \dots, P_k \in \mathbb{R}^n$ be polyhedra with $P = P_1 \cap \dots \cap P_k$, and let Q_i be an extended formulation for P_i for $i = 1, \dots, k$, i.e. $P_i = \{x \in \mathbb{R}^n : \exists y^{(i)} \in \mathbb{R}^{r_i} : (x, y^{(i)}) \in Q_i\}$. Then $P = \{x \in \mathbb{R}^n : \text{for } i = 1, \dots, k \exists y^{(i)} \in \mathbb{R}^{r_i} : (x, y^{(i)}) \in Q_i\}$.*

Theorem 5.14. *Let G be a perfect graph on n vertices. Then there is an algorithm that, on input G , produces an explicit extended formulation of $\text{STAB}(G)$ of size $n^{O(\log n)}$ in $n^{O(\log n)}$ time.*

Proof. We argue by induction on n . The base cases for n bounded by a constant are trivial, as the size of the classical formulation (and the time to obtain it) is constant too. For general n , Observation 5.13, together with Theorem 5.12, implies that we can obtain an extended formulation for $\text{STAB}(G)$ by adding together extended formulations of $\text{STAB}(G_0), \dots, \text{STAB}(G_k)$, where v_1, \dots, v_k are the vertices of G with degree at most $n/2$ and G_0, \dots, G_k are defined as above. First, assume that $k \geq n/2$, hence G_0, \dots, G_k have all size at most $n/2 + 1$. By induction, running the algorithm on G_0, \dots, G_k and then adding all the formulations obtained we get an extended formulation for $\text{STAB}(G)$ of size at most $n \cdot \left(\frac{n}{2} + 1\right)^{c \log(\frac{n}{2} + 1)}$ for some constant $c > 0$, but this is at most $n^{c \log n}$ (under the assumption, which can be made without loss of generality, that $c \geq 2$). The same bound holds for the total running time. Now if $k < n/2$, consider the complement graph \bar{G} , for which $k \geq n/2$, hence by the previous case the algorithm obtains a formulation of $\text{STAB}(\bar{G})$ of size at most $n \cdot \left(\frac{n}{2} + 1\right)^{c \log(\frac{n}{2} + 1)}$. We can then use Lemma 5.10 to efficiently obtain a formulation for $\text{STAB}(G)$, which by Corollary 5.11 has size at most

$n \cdot \left(\frac{n}{2} + 1\right)^{c \log(\frac{n}{2} + 1)} + 1 \leq n^{c \log n}$ (similar calculations work for the number of variables and equations of the formulation). Again, the same bound holds for the total running time. \square

Although in this section we restricted ourselves to perfect graphs for ease of exposition, we remark that the above algorithm can be used on general graphs, yielding an extended formulation of $(\text{STAB}(G), \text{QSTAB}(G))$. This can be easily seen by following similar arguments as the ones above, and in particular by noticing that the following holds for any graph G (where $k, G_0, \dots, G_k, V_0, \dots, V_k$ are defined as above):

- $\text{STAB}(G) \subseteq \{x \in \mathbb{R}_+^n : x^T y \leq 1 \ \forall y \in \text{QSTAB}(\tilde{G})\} \subseteq \text{QSTAB}(G)$;
- $\text{STAB}(G) \subseteq (\text{QSTAB}(G_1) \times \mathbb{R}^{V \setminus V_1}) \cap \dots \cap (\text{QSTAB}(G_k) \times \mathbb{R}^{V \setminus V_k}) = \text{QSTAB}(G)$.

One can see that the above inclusions are strict for non-perfect graphs, for instance for G equal to the cycle of length five.

We conclude by remarking that the formulation given by Theorem 5.14 can be seen as a more direct and slightly optimized version of the one given at the end of Section 5.3. Most notably, while the latter needs to take into account 1-rectangles as well as 0-rectangles, the former only explores the parts of the tree with lead to 1-rectangles, taking polars and intersections instead of convex hulls. This can be a significant advantage in practice.

5.4.3 Claw-free perfect graphs and generalizations

Let $P = \text{STAB}(G)$ where G is a claw-free, perfect graph on n vertices. As G is perfect, the (non-trivial part of the) slack matrix of P is the clique vs stable set incidence matrix of G , and can be computed by the following protocol, given in [30]. Alice, who has a clique C as input, sends a vertex $v \in C$ to Bob, who has a stable set S . Now, since G is claw-free, we have $|N(v) \cap S| \leq 2$, and clearly $C \subset N(v)$, hence Bob can just send $N(v) \cap S$ and Alice knows the intersection $C \cap S$. The protocol has complexity at most $3 \log n + 1$ hence by applying Theorem 5.2 we get the following formulation of size $O(n^3)$:

$$\begin{aligned} x(C) + \sum_{R \in \mathcal{R}_C} y_R &= 1 \quad \forall C \text{ clique of } G \\ y &\geq 0 \end{aligned} \tag{5.4}$$

Where \mathcal{R} contains a rectangle for each couple (v, U) , where $v \in V$ and $U \subset N(v)$ with U stable (i.e., $|U| \leq 2$), and \mathcal{R}_C , following the notation used in (5.2), denotes the set of rectangles including the row index corresponding to C . Notice that, for the rectangles in \mathcal{R} to partition the slack matrix of $\text{STAB}(G)$, we need to specify which vertex is sent from Alice given a certain

Chapter 5. Extended formulations in output-efficient time from communication protocols

clique as input: for this we can simply fix an order of the vertices and assume that Alice sends the vertex of her clique that is first in the order. Hence the rectangles in \mathcal{R}_C have form (v, U) where v is the “first” vertex of C . We now derive a more compact formulation than (5.4), getting rid of provably redundant equations. Before, we notice that the above protocol can be easily generalized to perfect $K_{1,t}$ -free graphs for $t \geq 3$: in this case the sets $\mathcal{R}, \mathcal{R}_C$ is defined similarly as before, except that now we have rectangles (v, U) with $|U| \leq t - 1$. This yields a formulation of size $O(n^t)$. We state our result for this more general class of graphs: informally, the only clique equations that we keep are coming from singletons and edges, obtaining a formulation with only $O(n^2)$ many equations.

Theorem 5.15. *Let $G(V, E)$ be a perfect and $K_{1,t}$ -free graph. Let $\mathcal{R}, \mathcal{R}_C$ as above. Then the following is an extended formulation for $\text{STAB}(G)$:*

$$\begin{aligned} x(v) + \sum_{R \in \mathcal{R}_v} y_R &= 1 \quad \forall v \in V \\ x(e) + \sum_{R \in \mathcal{R}_e} y_R &= 1 \quad \forall e \in E \\ y &\geq 0 \end{aligned} \tag{5.5}$$

Proof. Thanks to the above discussion, we only need to show that, for any clique C of G with $|C| = k \geq 3$, the equation $x(C) + \sum_{R \in \mathcal{R}_C} y_R = 1$ is implied by the equations in (5.5). From now on, fix such C and let $v \in C$ be the first vertex of C (in the order fixed by the protocol) and consider the following expression, obtained by summing the non-constant part of the equations relative to $e = uv$, for every $u \in C - v$:

$$\begin{aligned} \sum_{\substack{e=uv: \\ u \in C-v}} \left(x(e) + \sum_{R \in \mathcal{R}_e} y_R \right) &= \\ (k-2)x(v) + x(C) + \sum_{\substack{e=uv: \\ u \in C-v}} \sum_{R \in \mathcal{R}_e \cap \mathcal{R}_C} y_R + \sum_{R \in \mathcal{R}_e \setminus \mathcal{R}_C} y_R \end{aligned}$$

Now, recall the slack matrix of $\text{STAB}(G)$ has 0/1 entries and a 1-rectangle is determined by a couple (v, U) , where v is a vertex sent by Alice and U is the set of vertices sent by Bob. If the rectangle covers a 1-entry (C, S) , then v is the first vertex of C and $U = N(v) \cap S$, with $U \cap C = \emptyset$ (as otherwise (C, S) would be a 0-entry). Hence, we can derive $\mathcal{R}_e = \{(v, U) : U \subset N(v), U \in \mathcal{S}, u \notin U\}$ for $e = (u, v)$, and $\mathcal{R}_C = \{(v, U) : U \subset N(v), U \in \mathcal{S}, U \cap C = \emptyset\}$, where \mathcal{S} denotes the family of the stable sets of G . Hence $\mathcal{R}_C \subset \mathcal{R}_e$ for $e \in C$. We can then rewrite the above expression as:

$$\begin{aligned} (k-2)x(v) + x(C) + (k-1) \sum_{R \in \mathcal{R}_C} y_R + \sum_{\substack{e=uv: \\ u \in C}} \sum_{\substack{U \subset N(v)-u \\ U \in \mathcal{S}, U \cap C \neq \emptyset}} y_{v,U} &= \\ (k-2)x(v) + x(C) + (k-1) \sum_{R \in \mathcal{R}_C} y_R + (k-2) \sum_{\substack{U \subset N(v) \\ U \in \mathcal{S}, U \cap C \neq \emptyset}} y_{v,U}. \end{aligned} \tag{5.6}$$

Now, consider the right-hand side of the equation corresponding to $\{v\}$:

$$x(v) + \sum_{R \in \mathcal{R}_v} y_R = x(v) + \sum_{U \subset N(v), U \in \mathcal{S}} y_{v,U}.$$

Subtracting $k - 2$ times the latter from 5.6 we obtain $x(C) + \sum_{R \in \mathcal{R}_C} y_R$, which is the right-hand side that we wanted. Since we manipulated the non-constant part of equations whose space of solutions is non-empty, the corresponding constant part must be coherent (i.e. equal to 1) and we are done. \square

5.4.4 Comparability graphs

Let G be a comparability graph, and let (D, \leq_D) its underlying partial order. A clique (resp. stable set) in G corresponds to a chain (resp. antichain) in D . In [100], it is described an unambiguous nondeterministic protocol for the slack matrix of $\text{STAB}(G)$, which we now recall. Given a clique $C = \{v_1, \dots, v_k\}$ with $v_1 \leq \dots \leq v_k$ in D , and a stable set S disjoint from C , there are three cases: 1) every node of C precedes some node of S (equivalently, v_k does); 2) no node of C precedes a node of S (equivalently, v_1 does not precede any node of S); 3) there is an i such that v_i precedes some node of S , and v_{i+1} does not. Alice, given C , guesses which of the three cases applies and sends to Bob the certificate (v_k, L) (for last) in case 1), (v_1, F) (for first) in case 2) and (v_i, v_j) in case 3). This protocol yields a factorization of the slack matrix, hence an extended formulation for $\text{STAB}(G)$ of the usual kind:

$$\begin{aligned} x(C) + y(v_1, F) + y(v_1, v_2) + \dots + y(v_k, L) &= 1 \quad \forall C = \{v_1, \dots, v_k\} \in G \\ x, y &\geq 0 \end{aligned} \tag{5.7}$$

Lemma 5.16. *Let $G(V, E)$ be a comparability graph with order \leq_D , then the following is an extended formulation for $\text{STAB}(G)$:*

$$\begin{aligned} x(v) + y(v, F) + y(v, L) &= 1 \quad \forall v \in V \\ x(u) + x(v) + y(u, F) + y(u, v) + y(v, L) &= 1 \quad \forall u, v \in V : u \leq_D v \\ x, y &\geq 0. \end{aligned} \tag{5.8}$$

Proof. Let (x, y) be a point of (5.8), and $C = \{v_1, \dots, v_k\}$ a clique of G with $v_1 \leq_D \dots \leq_D v_k$, $k \geq 3$. Manipulating the equation of (5.8), we have that for $i = 2, \dots, k$, $x(v_i) = y(v_{i-1}, L) - y(v_{i-1}, v_i) - y(v_i, L)$. Hence:

$$\begin{aligned} x(C) &= \\ x(v_1) + y(v_1, L) - y(v_1, v_2) - y(v_2, L) + \dots + y(v_{k-1}, L) - y(v_{k-1}, v_k) - y(v_k, L) \\ &= x(v_1) + y(v_1, L) - y(v_1, v_2) - \dots - y(v_{k-1}, v_k) - y(v_k, L) \leq 1. \end{aligned}$$

\square

Chapter 5. Extended formulations in output-efficient time from communication protocols

Remark: An explicit extended formulation for the stable set polytopes of comparability graphs has been given in [74] by Lovasz. Both this formulation and the one given by us have quadratic size in the number of vertices of the graph, however Lovasz's formulation has only a linear number of variables.

6 Conclusion

In this thesis we examined the subject of 2-level polytopes from various perspectives. Our purpose was to draw attention to interesting open problems that are, in our opinion, not studied enough, and of course to describe the progress we made toward their solution. We would now like to conclude this thesis by pointing out the main research directions that stem from our work.

The reason we first started to work on 2-level polytopes was our interest in Conjecture 2.1. Even though we succeeded to prove it for essentially all combinatorial classes we could identify, the general case remains open: what bound can we prove on the product of the number of vertices and facets of any 2-level polytope? The only bound that we know is 4^d for dimension d , and, as argued in Section ??, this is not tight. This is not only a problem of intrinsic interest, but it might be the most approachable open question on general 2-level polytopes, hopefully paving the way for the harder questions we are now going to describe.

The question on the extension complexity of 2-level polytopes is probably the most meaningful and fascinating that we approached, for the various connections already described in the introduction. In Chapter 3, Theorem 3.11 we give the first and only known non-trivial lower bound on the extension complexity of a 2-level polytope, in particular the stable set polytope of some bipartite graphs. While this only improves by a logarithmic factor on the trivial bound, we hope that the same or a similar technique can be applied to other graphs to obtain stronger bounds. Of course other classes of 2-level polytopes might be more promising for this purpose, especially if one aims at proving superpolynomial lower bounds. However it is worth to notice that there is no clear candidate for this task. A look at the various classes of 2-level polytopes studied in Chapter 2 reveals that most of them have a polynomial number of facets. An exception is the class of min up/down polytopes (see Section 2.3.3 or [72]): however it is known and not too hard to see (but not published as far as we know) that such polytopes have polynomial extension complexity. Hence so far we are not aware of any 2-level polytope that comes from a combinatorial context and could have high extension complexity (apart from the stable set polytopes of perfect graphs). It seems that to prove a strong lower bound one would have to resort to less structured polytopes, arising for instance from slices of the unit cube or from the “hypergraph embedding” given in [10], which has a combinatorial flavour

although it is general enough to describe *all* 2-level polytopes. On the other hand, it might very well be that a subexponential upper bound holds for the extension complexity of all 2-level polytopes. This would generalize Yannakakis' quasipolynomial bound for stable set polytopes of perfect graphs (Theorem 5.4). This direction has been considered by researchers who studied 2-level polytopes, but the current understanding of the subject still seems too poor: it is not clear at all how, for instance, a generalization of Yannakakis' protocol could be applied to all 2-level polytopes.

A different direction that might contribute to the above question, while being interesting on its own, stems from the problem of efficiently recognizing 0/1 slack matrices, which is studied in Chapter 4. In particular we feel that our decomposition approach via the operation of k -sum deserves further investigation. While we could only prove its successful application for the special case of 2-level matroid polytopes, the operation of k -sum might prove useful in more general contexts. For a moment, while working on the main results of the chapter, we considered the following bold conjecture: every 2-level polytope is obtained via k -sum from lower dimensional 2-level polytopes, or belongs to one of a few “basic” classes, which are “simple” and relatively well understood. This is inspired by the numerous decomposition theorems for perfect graphs [15] (while their stable set polytopes might be considered as one of the basic classes). Notice that proving such a result would have at least two relevant consequences: in light of Corollary 4.12, it would imply a bound on the extension complexity of 2-level polytopes; thanks to Theorem 4.25, it would imply an efficient algorithm for recognizing 0/1 slack matrices. Of course, as the numerical experiments in Section 4.4.3 show, such a statement seems to be false: there are many 2-level polytopes that are not k -sums, and it seems unlikely that they all belong to some special class. However, the same data suggests that k -sums do play a significant role in the context of 2-level polytopes, hence giving hope that introducing some new, maybe more complex operation might finally lead to a decomposition result. This would dramatically improve our understanding of the structure of 2-level polytopes, and might settle the most important questions we have on them.

Finally, Chapter 5 might leave the reader to wonder about many questions, possibly more than the rest of the thesis. We would first like to point out that, although we only apply our results to stable set polytopes of perfect graphs, Theorem 5.6 lends itself to applications well beyond the realm of 2-level polytopes. In particular it could be applied to non-2-level stable set polytopes. Recall that the algorithms described in Sections 5.3.1 and 5.4.2 give quasipolynomial extended formulations for $(\text{STAB}(G), \text{QSTAB}(G))$, for any graph G . How strong are these relaxations for non-perfect graphs? Interestingly, both $\text{STAB}(G)$ and $\text{QSTAB}(G)$ are in general NP-hard to optimize over, while our formulations are of quasipolynomial size in the worst case.

One last issue that is worth mentioning concerns the relationship between deterministic communication complexity and 2-level polytopes. As outlined in Chapter 1, the log-rank conjecture is a fundamental open problem that concerns the deterministic communication complexity of boolean matrices [75]. Since 2-level polytopes have 0/1 slack matrices (Observation 1.4), the bound of the conjecture would imply a $2^{\text{polylog}(d)}$ bound on the extension

complexity of d -dimensional 2-level polytopes. Moreover, such a bound would come from a deterministic protocol. The best partial progress on the conjecture is due to Lovett [75], and it implies a bound of $2^{O(\sqrt{d})}$. This is the best upper bound currently known for 2-level polytopes, suggesting that (deterministic) communication complexity might be the right perspective to approach the problem. As a side notice we mention that, thanks to Theorem 5.6, this bound might in principle be turned into an explicit extended formulation for all 2-level polytopes, but Lovett’s approach seems inherently non-constructive, not giving any explicit protocol to start with. Overall, it seems natural to ask what is really the role of deterministic protocols in the context of boolean slack matrices. In [30] it is shown that no deterministic protocol can yield a polynomial extended formulation for the spanning tree polytope, while a simple randomized protocol yields a cubic size formulation similar to Martin’s formulation [79]. Is there a 2-level polytope exhibiting a similar gap, or are deterministic protocols as powerful as randomized ones when it comes to boolean matrices? We believe that this and similar questions are worth asking in order to improve our understanding on the subject, and that deep answers wait to be brought to light.

A

A.1 The polytopes from Proposition 2.11, Chapter 2

The following are the `polymake` vertex descriptions of the two 8-dimensional polytopes from Proposition 2.11: the min up/down polytope $P_8(2)$ is denoted by `$P`, and the Hansen polytope $\text{Hans}(P_7)$ of the path on 7 nodes P_7 is denoted by `$H`.

```
$P = new Polytope(VERTICES=> [
[1, 0, 0, 0, 0, 0, 0, 0, 0, 0],[1, 1, 1, 1, 1, 1, 1, 1, 1, 1],
[1, 0, 0, 0, 0, 0, 0, 0, 0, 1],[1, 1, 1, 1, 1, 1, 1, 1, 1, 0],
[1, 0, 0, 0, 0, 0, 0, 0, 1, 1],[1, 1, 1, 1, 1, 1, 1, 1, 0, 0],
[1, 0, 0, 0, 0, 0, 0, 1, 1, 1],[1, 1, 1, 1, 1, 1, 0, 0, 0, 0],
[1, 0, 0, 0, 0, 0, 0, 1, 1, 0],[1, 1, 1, 1, 1, 1, 0, 0, 0, 1],
[1, 0, 0, 0, 0, 0, 1, 1, 1, 1],[1, 1, 1, 1, 1, 0, 0, 0, 0, 0],
[1, 0, 0, 0, 0, 0, 1, 1, 1, 0],[1, 1, 1, 1, 1, 0, 0, 0, 0, 1],
[1, 0, 0, 0, 0, 0, 1, 1, 0, 0],[1, 1, 1, 1, 1, 0, 0, 0, 1, 1],
[1, 0, 0, 0, 0, 1, 1, 1, 1, 1],[1, 1, 1, 1, 0, 0, 0, 0, 0, 0],
[1, 0, 0, 0, 0, 1, 1, 1, 1, 0],[1, 1, 1, 1, 0, 0, 0, 0, 0, 1],
[1, 0, 0, 0, 0, 1, 1, 0, 0, 0],[1, 1, 1, 1, 0, 0, 0, 1, 1, 1],
[1, 0, 0, 0, 0, 1, 1, 0, 0, 1],[1, 1, 1, 1, 0, 0, 0, 1, 1, 0],
[1, 0, 0, 0, 0, 1, 1, 0, 0, 0],[1, 1, 1, 1, 0, 0, 0, 1, 1, 1],
[1, 0, 0, 0, 1, 1, 1, 1, 1, 1],[1, 1, 1, 0, 0, 0, 0, 0, 0, 0],
[1, 0, 0, 0, 1, 1, 1, 1, 1, 0],[1, 1, 1, 0, 0, 0, 0, 0, 0, 1],
[1, 0, 0, 0, 1, 1, 1, 0, 0, 0],[1, 1, 1, 0, 0, 0, 0, 1, 1, 1],
[1, 0, 0, 0, 1, 1, 0, 0, 0, 0],[1, 1, 1, 0, 0, 0, 1, 1, 1, 0],
[1, 0, 0, 0, 1, 1, 0, 0, 0, 1],[1, 1, 1, 0, 0, 0, 1, 1, 0, 0],
[1, 0, 1, 1, 1, 1, 1, 1, 1, 1],[1, 1, 0, 0, 0, 0, 0, 0, 0, 0],
[1, 0, 1, 1, 1, 1, 1, 1, 1, 0],[1, 1, 0, 0, 0, 0, 0, 0, 0, 1]
```

Appendix A. An appendix

```
[1, 0, 1, 1, 1, 1, 1, 0, 0],[1, 1, 0, 0, 0, 0, 0, 1, 1],
[1, 0, 1, 1, 1, 1, 0, 0, 0],[1, 1, 0, 0, 0, 0, 1, 1, 1],
[1, 0, 1, 1, 1, 1, 0, 0, 1],[1, 1, 0, 0, 0, 0, 1, 1, 0],
[1, 0, 1, 1, 1, 0, 0, 0, 0],[1, 1, 0, 0, 0, 1, 1, 1, 1],
[1, 0, 1, 1, 1, 0, 0, 0, 1],[1, 1, 0, 0, 0, 1, 1, 1, 0],
[1, 0, 1, 1, 1, 0, 0, 1, 1],[1, 1, 0, 0, 0, 1, 1, 0, 0],
[1, 0, 1, 1, 0, 0, 0, 0, 0],[1, 1, 0, 0, 1, 1, 1, 1, 1],
[1, 0, 1, 1, 0, 0, 0, 0, 1],[1, 1, 0, 0, 1, 1, 1, 1, 0],
[1, 0, 1, 1, 0, 0, 0, 1, 1],[1, 1, 0, 0, 1, 1, 1, 0, 0],
[1, 0, 1, 1, 0, 0, 1, 1, 1],[1, 1, 0, 0, 1, 1, 0, 0, 0],
[1, 0, 1, 1, 0, 0, 1, 1, 0],[1, 1, 0, 0, 1, 1, 0, 0, 1]] );

$H= new Polytope(VERTICES=> [
[1, 1, 0, 0, 0, 0, 0, 0, 0],[1, -1, 0, 0, 0, 0, 0, 0, 0],
[1, 1, 0, 0, 0, 0, 0, 0, 1],[1, -1, 0, 0, 0, 0, 0, 0, -1],
[1, 1, 0, 0, 0, 0, 0, 1, 0],[1, -1, 0, 0, 0, 0, 0, -1, 0],
[1, 1, 0, 0, 0, 0, 1, 0, 0],[1, -1, 0, 0, 0, 0, -1, 0, 0],
[1, 1, 0, 0, 0, 0, 1, 0, 1],[1, -1, 0, 0, 0, 0, -1, 0, -1],
[1, 1, 0, 0, 0, 1, 0, 0, 0],[1, -1, 0, 0, 0, -1, 0, 0, 0],
[1, 1, 0, 0, 0, 1, 0, 0, 1],[1, -1, 0, 0, 0, -1, 0, 0, -1],
[1, 1, 0, 0, 0, 1, 0, 1, 0],[1, -1, 0, 0, 0, -1, 0, -1, 0],
[1, 1, 0, 0, 1, 0, 0, 0, 0],[1, -1, 0, 0, -1, 0, 0, 0, 0],
[1, 1, 0, 0, 1, 0, 0, 0, 1],[1, -1, 0, 0, -1, 0, 0, 0, -1],
[1, 1, 0, 0, 1, 0, 0, 1, 0],[1, -1, 0, 0, -1, 0, 0, -1, 0],
[1, 1, 0, 0, 1, 0, 1, 0, 0],[1, -1, 0, 0, -1, 0, -1, 0, 0],
[1, 1, 0, 0, 1, 0, 1, 0, 1],[1, -1, 0, 0, -1, 0, -1, 0, -1],
[1, 1, 0, 1, 0, 0, 0, 0, 0],[1, -1, 0, -1, 0, 0, 0, 0, 0],
[1, 1, 0, 1, 0, 0, 0, 0, 1],[1, -1, 0, -1, 0, 0, 0, 0, -1],
[1, 1, 0, 1, 0, 0, 0, 1, 0],[1, -1, 0, -1, 0, 0, 0, -1, 0],
[1, 1, 0, 1, 0, 0, 1, 0, 0],[1, -1, 0, -1, 0, 0, -1, 0, 0],
[1, 1, 0, 1, 0, 0, 1, 0, 1],[1, -1, 0, -1, 0, 0, -1, 0, -1],
[1, 1, 0, 1, 0, 1, 0, 0, 0],[1, -1, 0, -1, 0, -1, 0, 0, 0],
[1, 1, 0, 1, 0, 1, 0, 0, 1],[1, -1, 0, -1, 0, -1, 0, 0, -1],
[1, 1, 1, 0, 0, 0, 0, 0, 0],[1, -1, -1, 0, 0, 0, 0, 0, 0],
[1, 1, 1, 0, 0, 0, 0, 0, 1],[1, -1, -1, 0, 0, 0, 0, 0, -1],
[1, 1, 1, 0, 0, 0, 0, 1, 0],[1, -1, -1, 0, 0, 0, 0, -1, 0],
[1, 1, 1, 0, 0, 0, 1, 0, 0],[1, -1, -1, 0, 0, 0, -1, 0, 0],
[1, 1, 1, 0, 0, 0, 1, 0, 1],[1, -1, -1, 0, 0, 0, -1, 0, -1],
[1, 1, 1, 0, 0, 1, 0, 0, 0],[1, -1, -1, 0, 0, -1, 0, 0, 0],
[1, 1, 1, 0, 0, 1, 0, 0, 1],[1, -1, -1, 0, 0, -1, 0, 0, -1],
[1, 1, 1, 0, 0, 1, 0, 1, 0],[1, -1, -1, 0, 0, -1, 0, -1, 0],
```

```
[1, 1, 1, 0, 1, 0, 0, 0, 0, 0],[1, -1, -1, 0, -1, 0, 0, 0, 0, 0],
[1, 1, 1, 0, 1, 0, 0, 0, 1, 1],[1, -1, -1, 0, -1, 0, 0, 0, 0, -1],
[1, 1, 1, 0, 1, 0, 0, 1, 0, 0],[1, -1, -1, 0, -1, 0, 0, -1, 0, 0],
[1, 1, 1, 0, 1, 0, 1, 0, 0, 0],[1, -1, -1, 0, -1, 0, -1, 0, 0, 0],
[1, 1, 1, 0, 1, 0, 1, 0, 1, 1],[1, -1, -1, 0, -1, 0, -1, 0, -1, 0]] );
```

A.2 The polytope from Section 2.6.5

The following is the polymake inequality description for the 12-dimensional polytope from Example 2.6.5.

```
$c= new Polytope(INEQUALITIES=>[
[0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 0, 1, 1] ,
[1, -1, -1, -1, -1, -1, -1, -1, -1, -1, 0, -1, -1] ,
[0, 1, 0, 1, 0, 1, 1, 1, 0, 1, 1, 0, 1] ,
[1, -1, 0, -1, 0, -1, -1, -1, 0, -1, -1, 0, -1] ,
[0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 1, 1, 0, 1] ,
[1, 0, 0, 0, 0, -1, 0, 0, 0, 0, -1, -1, 0, -1] ,
[0, 1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0] ,
[1, -1, 0, 0, 0, 0, 0, 0, 0, -1, 0, 0, 0, 0] ,
[0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1] ,
[1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1] ,
[0, 1, 1, 1, 1, 1, 0, 0, 0, 0, 0, 0, 1, 1] ,
[1, -1, -1, -1, -1, -1, 0, 0, 0, 0, 0, 0, -1, -1] ,
[0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0] ,
[1, -1, 0, 0, 0, 0, 0, 0, 0, 0, 0, -1, 0, 0] ,
[0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 1] ,
[1, -1, 0, 0, 0, 0, 0, 0, 0, -1, 0, 0, 0, -1] ,
[0, 1, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0] ,
[1, -1, -1, 0, -1, 0, 0, 0, 0, 0, 0, -1, 0, 0] ,
[0, 1, 1, 1, 0, 1, 1, 0, 1, 1, 0, 1, 0, 1] ,
[1, -1, -1, -1, 0, -1, -1, 0, -1, -1, 0, -1, 0, 1] ,
[0, 0, 1, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1] ,
[1, 0, -1, 0, 0, 0, 0, -1, 0, 0, 0, 0, 0, -1] ,
[0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 0, 1, 1] ,
[1, -1, -1, -1, -1, -1, -1, -1, -1, -1, 0, -1, -1] ,
[0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0] ,
[1, -1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0] ,
[0, 1, 1, 1, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0] ,
[1, -1, -1, -1, 0, 0, 0, 0, 0, -1, -1, 0, 0, 0] ,
[0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1] ,
[1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1] ,
```

Appendix A. An appendix

[0, 1, 1, 1, 0, 0, 1, 1, 1, 1, 0, 1, 1] ,
[1, -1, -1, -1, 0, 0, -1, -1, -1, -1, 0, -1, -1]]);

Bibliography

- [1] M. Aprile, A. Cevallos, and Y. Faenza. “On 2-level polytopes arising in combinatorial settings”. In: *SIAM Journal on Discrete Mathematics* 32.3 (2018), pp. 1857–1886.
- [2] M. Aprile, A. Cevallos, and Y. Faenza. “On vertices and facets of combinatorial 2-level polytopes”. In: *International Symposium on Combinatorial Optimization*. Springer, 2016, pp. 177–188.
- [3] M. F. Aprile, Y. Faenza, S. Fiorini, T. Huynh, and M. Macchia. “Extension complexity of stable set polytopes of bipartite graphs”. In: *Lecture Notes in Computer Science book series (LNCS)* 10520.CONF (2017), pp. 75–87.
- [4] E. Balas. “Disjunctive programming”. In: *Annals of Discrete Mathematics*. Vol. 5. Elsevier, 1979, pp. 3–51.
- [5] F. Barahona and M. Grötschel. “On the cycle polytope of a binary matroid”. In: *Journal of combinatorial theory, Series B* 40.1 (1986), pp. 40–62.
- [6] F. Barahona and A. R. Mahjoub. “On the cut polytope”. In: *Mathematical programming* 36.2 (1986), pp. 157–173.
- [7] I. Bárány and A. Pór. “On 0-1 polytopes with many facets”. In: *Advances in Mathematics* 161.2 (2001), pp. 209–228.
- [8] C. Berge. “Some classes of perfect graphs”. In: *Six Papers on Graph Theory [related to a series of lectures at the Research and Training School of the Indian Statistical Institute, Calcutta]* (1963), pp. 1–21.
- [9] R. E. Bixby and W. H. Cunningham. *Matroid optimization and algorithms*. MIT Press, 1996.
- [10] A. Bohn, Y. Faenza, S. Fiorini, V. Fisikopoulos, M. Macchia, and K. Pashkovich. “Enumeration of 2-Level Polytopes”. In: *Algorithms – ESA 2015*. Vol. 9294. Lecture Notes in Computer Science, 2015, pp. 191–202. ISBN: 978-3-662-48349-7. DOI: 10.1007/978-3-662-48350-3_17. URL: http://dx.doi.org/10.1007/978-3-662-48350-3_17.
- [11] A. Bohn, Y. Faenza, S. Fiorini, V. Fisikopoulos, M. Macchia, and K. Pashkovich. “Enumeration of 2-level polytopes”. In: *arXiv preprint arXiv:1703.01943* (2017).
- [12] C. Bron and J. Kerbosch. “Algorithm 457: finding all cliques of an undirected graph”. In: *Communications of the ACM* 16.9 (1973), pp. 575–577.

Bibliography

- [13] F. Buekenhout and M. Parker. “The number of nets of the regular convex polytopes in dimension ≤ 4 ”. In: *Discrete mathematics* 186.1 (1998), pp. 69–94.
- [14] T. Chappell, T. Friedl, and R. Sanyal. “Two double poset polytopes”. In: *SIAM Journal on Discrete Mathematics* 31.4 (2017), pp. 2378–2413.
- [15] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas. “The strong perfect graph theorem”. In: *Annals of mathematics* (2006), pp. 51–229.
- [16] V. Chvátal. “On certain polytopes associated with graphs”. In: *J. Combinatorial Theory Ser. B* 18 (1975), pp. 138–154.
- [17] V. Chvátal. “Perfectly ordered graphs”. In: *North-Holland mathematics studies* 88 (1984), pp. 63–65.
- [18] V. Chvátal. “On certain polytopes associated with graphs”. In: *Journal of Combinatorial Theory, Series B* 18.2 (1975), pp. 138–154.
- [19] M. Conforti, V. Kaibel, M. Walter, and S. Weltge. “Subgraph polytopes and independence polytopes of count matroids”. In: *Operations Research Letters* 43.5 (2015), pp. 457–460.
- [20] M. Conforti, G. Cornuéjols, and G. Zambelli. “Extended formulations in combinatorial optimization”. In: *Ann. Oper. Res.* 204 (2013), pp. 97–143. ISSN: 0254-5330. DOI: 10.1007/s10479-012-1269-0. URL: <http://dx.doi.org/10.1007/s10479-012-1269-0>.
- [21] M. Conforti, G. Cornuéjols, and G. Zambelli. *Integer programming*. Vol. 271. Graduate Texts in Mathematics. Springer, Cham, 2014, pp. xii+456. ISBN: 978-3-319-11007-3; 978-3-319-11008-0. DOI: 10.1007/978-3-319-11008-0. URL: <http://dx.doi.org/10.1007/978-3-319-11008-0>.
- [22] M. Conforti, B. Gerards, and K. Pashkovich. “Stable sets and graphs with no even holes”. In: *Mathematical Programming* 153.1 (2015), pp. 13–39.
- [23] G. Cornuéjols, X. Liu, and K. Vuskovic. “A polynomial algorithm for recognizing perfect graphs”. In: *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*. IEEE, 2003, pp. 20–27.
- [24] T. M. Cover and J. A. Thomas. “Elements of information theory 2nd edition”. In: (2006).
- [25] H. S. M. Coxeter. *Projective geometry*. Revised reprint of the second (1974) edition. Springer-Verlag, New York, 1994, pp. xii+162. ISBN: 0-387-96532-7.
- [26] J. Edmonds. “Matroids and the greedy algorithm”. In: *Math. Programming* 1 (1971), pp. 127–136. ISSN: 0025-5610. DOI: 10.1007/BF01584082. URL: <http://dx.doi.org/10.1007/BF01584082>.
- [27] J. Edmonds. “Paths, trees, and flowers”. In: *Canad. J. Math.* 17 (1965), pp. 449–467. ISSN: 0008-414X. DOI: 10.4153/CJM-1965-045-4. URL: <http://dx.doi.org/10.4153/CJM-1965-045-4>.
- [28] P. Eirinakis, D. Magos, I. Mourtos, and P. Miliotis. “Polyhedral aspects of stable marriage”. In: *Mathematics of Operations Research* 39.3 (2013), pp. 656–671.

- [29] M. Escobar and Y. Faenza. *Rotation-perfect stable marriage instances*. Manuscript. 2017.
- [30] Y. Faenza, S. Fiorini, R. Grappe, and H. R. Tiwary. “Extended formulations, nonnegative factorizations, and randomized communication protocols”. In: *Mathematical Programming* 153.1 (2015), pp. 75–94.
- [31] E. M. Feichtner and B. Sturmfels. “Matroid polytopes, nested sets and Bergman fans”. In: *Portugaliae Mathematica* 62.4 (2005), pp. 437–468.
- [32] S. Fiorini, T. Huynh, and S. Weltge. “Strengthening Convex Relaxations of 0/1-Sets Using Boolean Formulas”. In: *arXiv preprint arXiv:1711.01358* (2017).
- [33] S. Fiorini, V. Kaibel, K. Pashkovich, and D. O. Theis. “Combinatorial bounds on nonnegative rank and extended formulations”. In: *Discrete Math.* 313.1 (2013), pp. 67–83. ISSN: 0012-365X. DOI: 10.1016/j.disc.2012.09.015. URL: <http://dx.doi.org/10.1016/j.disc.2012.09.015>.
- [34] S. Fiorini, M. Macchia, and K. Pashkovich. “Bounds on the number of 2-level polytopes, cones and configurations”. In: *arXiv preprint arXiv:1806.06011* (2018).
- [35] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. De Wolf. “Linear vs. semidefinite extended formulations: exponential separation and strong lower bounds”. In: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM. 2012, pp. 95–106.
- [36] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf. “Exponential lower bounds for polytopes in combinatorial optimization”. In: *J. ACM* 62.2 (2015), Art. 17, 23. ISSN: 0004-5411. DOI: 10.1145/2716307. URL: <http://dx.doi.org/10.1145/2716307>.
- [37] S. Fujishige. “A characterization of faces of the base polyhedron associated with a submodular system”. In: *Journal of the Operations Research Society of Japan* 27.2 (1984), pp. 112–129.
- [38] D. Gale and L. S. Shapley. “College admissions and the stability of marriage”. In: *The American Mathematical Monthly* 69.1 (1962), pp. 9–15.
- [39] I. M. Gelfand, R. M. Goresky, R. D. MacPherson, and V. V. Serganova. “Combinatorial geometries, convex polyhedra, and Schubert cells”. In: *Advances in Mathematics* 63.3 (1987), pp. 301–316.
- [40] I. Gleason and I. Hubbard. “Products of abstract polytopes”. In: *arXiv preprint arXiv:1603.03585* (2016).
- [41] M. Göös, R. Jain, and T. Watson. “Extension complexity of independent set polytopes”. In: *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*. IEEE. 2016, pp. 565–572.
- [42] M. Göös, R. Jain, and T. Watson. “Extension complexity of independent set polytopes”. In: *SIAM Journal on Computing* 47.1 (2018), pp. 241–269.

Bibliography

- [43] J. Gouveia, R. Grappe, V. Kaibel, K. Pashkovich, R. Robinson, and R. Thomas. “Which nonnegative matrices are slack matrices?” In: *Linear Algebra and its Applications* 439.10 (2013), pp. 2921–2933.
- [44] J. Gouveia, M. Laurent, P. A. Parrilo, and R. Thomas. “A new semidefinite programming hierarchy for cycles in binary matroids and cuts in graphs”. In: *Math. programming* 133.1-2 (2012), pp. 203–225.
- [45] J. Gouveia, P. Parrilo, and R. Thomas. “Theta bodies for polynomial ideals”. In: *SIAM Journal on Optimization* 20.4 (2010), pp. 2097–2118.
- [46] J. Gouveia, K. Pashkovich, R. Z. Robinson, and R. R. Thomas. “Four-dimensional polytopes of minimum positive semidefinite rank”. In: *J. Comb. Theory, Ser. A* 145 (2017), pp. 184–226. DOI: 10.1016/j.jcta.2016.08.002. URL: <https://doi.org/10.1016/j.jcta.2016.08.002>.
- [47] J. Gouveia, R. Z. Robinson, and R. R. Thomas. “Polytopes of Minimum Positive Semidefinite Rank”. In: *Discrete & Computational Geometry* 50.3 (2013), pp. 679–699. DOI: 10.1007/s00454-013-9533-x. URL: <https://doi.org/10.1007/s00454-013-9533-x>.
- [48] F. Grande and R. Sanyal. “Theta rank, levelness, and matroid minors”. In: *Journal of Combinatorial Theory, Series B* 123 (2017), pp. 1–31.
- [49] J. L. Gross and J. Yellen. *Graph theory and its applications*. CRC press, 2005.
- [50] M. Grötschel, L. Lovász, and A. Schrijver. “The ellipsoid method and its consequences in combinatorial optimization”. In: *Combinatorica* 1.2 (1981), pp. 169–197.
- [51] D. Gusfield. “Three fast algorithms for four problems in stable marriage”. In: *SIAM Journal on Computing* 16.1 (1987), pp. 111–128.
- [52] D. Gusfield and R. W. Irving. *The stable marriage problem: structure and algorithms*. MIT press, 1989.
- [53] O. Hanner. “Intersections of translates of convex bodies”. In: *Mathematica Scandinavica* 4 (1956), pp. 65–87.
- [54] A. Hansen. “On a certain class of polytopes associated with independence systems.” In: *Mathematica Scandinavica* 41 (1977), pp. 225–241.
- [55] T. Hibi and N. Li. “Unimodular equivalence of order and chain polytopes”. In: *arXiv preprint arXiv:1208.4029* (2012).
- [56] R. W. Irving and P. Leather. “The complexity of counting stable marriages”. In: *SIAM Journal on Computing* 15.3 (1986), pp. 655–667.
- [57] S. Iwata, N. Kamiyama, N. Katoh, S. Kijima, and Y. Okamoto. “Extended formulations for sparsity matroids”. In: *Mathematical Programming* 158.1-2 (2016), pp. 565–574.
- [58] V. Kaibel, J. Lee, M. Walter, and S. Weltge. “Extended formulations for independence polytopes of regular matroids”. In: *Graphs and Combinatorics* 32.5 (2016), pp. 1931–1944.

- [59] V. Kaibel. “Extended formulations in combinatorial optimization”. In: *arXiv preprint arXiv:1104.1023* (2011).
- [60] V. Kaibel and G. Averkov. *Personal communication*.
- [61] V. Kaibel and M. E. Pfetsch. “Some algorithmic problems in polytope theory”. In: *Algebra, geometry and software systems*. Springer, 2003, pp. 23–47.
- [62] G. Kalai. “The number of faces of centrally-symmetric polytopes”. In: *Graphs and Combinatorics* 5.1 (1989), pp. 389–391. ISSN: 0911-0119.
- [63] L. Khachiyan, E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. “On the complexity of some enumeration problems for matroids”. In: *SIAM Journal on Discrete Mathematics* 19.4 (2005), pp. 966–984.
- [64] K. Khoshkhah and D. O. Theis. “Fooling Sets and the Spanning Tree Polytope”. In: *arXiv preprint arXiv:1701.00350* (2017).
- [65] K. Khoshkhah and D. O. Theis. “On the Combinatorial Lower Bound for the Extension Complexity of the Spanning Tree Polytope”. In: *arXiv preprint arXiv:1702.01424* (2017).
- [66] D. E. Knuth. *Mariages stables et leurs relations avec d'autres problèmes combinatoires: introduction à l'analyse mathématique des algorithmes*. Montréal: Presses de l'Université de Montréal, 1976.
- [67] P. Kolman, M. Koutecký, and H. R. Tiwary. “Extension Complexity, MSO Logic, and Treewidth”. In: *LIPICs-Leibniz International Proceedings in Informatics*. Vol. 53. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- [68] U. Kortenkamp, J. Richter-Gebert, A. Sarangarajan, and G. M. Ziegler. “Extremal Properties of 0/1-Polytopes”. In: *Discrete & Computational Geometry* 17.4 (1997), pp. 439–448. DOI: 10.1007/PL00009303. URL: <https://doi.org/10.1007/PL00009303>.
- [69] K. Krishnan and J. E. Mitchell. “A linear programming approach to semidefinite programming problems”. PhD thesis. PhD thesis, Mathematical Sciences, Rensselaer Polytechnic Institute, 2002.
- [70] E. Kushilevitz and N. Nisan. “Communication Complexity”. In: (1996).
- [71] M. Laurent. “A comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre relaxations for 0–1 programming”. In: *Mathematics of Operations Research* 28.3 (2003), pp. 470–496.
- [72] J. Lee, J. Leung, and F. Margot. “Min-up/min-down polytopes”. In: *Discrete Optimization* 1.1 (2004), pp. 77–85.
- [73] L. Lovász. “On the Shannon capacity of a graph”. In: *IEEE Transactions on Information theory* 25.1 (1979), pp. 1–7.
- [74] L. Lovász. “Stable sets and polynomials”. In: *Discrete mathematics* 124.1-3 (1994), pp. 137–153.
- [75] S. Lovett. “Communication is bounded by root of rank”. In: *Journal of the ACM (JACM)* 63.1 (2016), p. 1.

Bibliography

- [76] M. Macchia. “Two level polytopes: geometry and optimization”. PhD thesis. Université libre de Bruxelles, 2018.
- [77] E. Margot. *Composition de polytopes combinatoires: une approche par projection*. Vol. 4. PPUR presses polytechniques, 1995.
- [78] R. K. Martin. “Using separation algorithms to generate mixed integer model reformulations”. In: *Operations Research Letters* 10.3 (1991), pp. 119–128.
- [79] R. K. Martin. “Using separation algorithms to generate mixed integer model reformulations”. In: *Oper. Res. Lett.* 10.3 (1991), pp. 119–128. ISSN: 0167-6377. DOI: 10.1016/0167-6377(91)90028-N. URL: [http://dx.doi.org/10.1016/0167-6377\(91\)90028-N](http://dx.doi.org/10.1016/0167-6377(91)90028-N).
- [80] J. G. Oxley. *Matroid theory*. Vol. 3. Oxford University Press, USA, 2006.
- [81] K. Pashkovich. “Extended formulations for combinatorial polytopes”. PhD thesis. Otto-von-Guericke-Universität Magdeburg, 2012.
- [82] H. Prodinger and R. F. Tichy. “Fibonacci numbers of graphs”. In: *Fibonacci Quarterly* 20.1 (1982), pp. 16–21.
- [83] M. Queyranne. “Minimizing symmetric submodular functions”. In: *Mathematical Programming* 82.1-2 (1998), pp. 3–12.
- [84] A. E. Roth, U. G. Rothblum, and J. H. Vande Vate. “Stable matchings, optimal assignments, and linear programming”. In: *Mathematics of operations research* 18.4 (1993), pp. 803–828.
- [85] T. Rothvoß. “Some 0/1 polytopes need exponential size extended formulations”. In: *Mathematical Programming* 142.1-2 (2013), pp. 255–268.
- [86] T. Rothvoß. “The matching polytope has exponential extension complexity”. In: *STOC’14—Proceedings of the 2014 ACM Symposium on Theory of Computing*. ACM, New York, 2014, pp. 263–272.
- [87] T. Roughgarden. “Communication complexity (for algorithm designers)”. In: *arXiv preprint arXiv:1509.06257* (2015).
- [88] F. Santos. “A counterexample to the Hirsch Conjecture”. In: *Annals of Mathematics* 176.1 (2012), pp. 383–412.
- [89] A. Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Vol. 24. Springer Science & Business Media, 2002.
- [90] A. Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Vol. 24. Springer Science & Business Media, 2003.
- [91] A. Schrijver. *Combinatorial Optimization. Polyhedra and Efficiency*. Springer, 2003.
- [92] R. Stanley. “Two poset polytopes”. In: *Discrete & Computational Geometry* 1.1 (1986), pp. 9–23.
- [93] R. P. Stanley. “Decompositions of rational convex polytopes”. In: *Ann. Discrete Math* 6.6 (1980), pp. 333–342.

-
- [94] S. Sullivant. “Compressed polytopes and statistical disclosure limitation”. In: *Tohoku Mathematical Journal, Second Series* 58.3 (2006), pp. 433–445.
 - [95] W. T. Tutte. “Lectures on matroids”. In: *J. Res. Nat. Bur. Standards Sect. B* 69.1-47 (1965), p. 468.
 - [96] Z. Tuza. “Covering of graphs by complete bipartite subgraphs: complexity of 0-1 matrices”. In: *Combinatorica* 4.1 (1984), pp. 111–116. ISSN: 0209-9683. DOI: 10.1007/BF02579163. URL: <http://dx.doi.org/10.1007/BF02579163>.
 - [97] S. Wagner and I. Gutman. “Maxima and minima of the Hosoya index and the Merrifield-Simmons index”. In: *Acta Applicandae Mathematicae* 112.3 (2010), pp. 323–346.
 - [98] S. Weltge. “Sizes of linear descriptions in combinatorial optimization”. PhD thesis. Otto-von-Guericke-Universität Magdeburg, Fakultät für Mathematik, 2015.
 - [99] R. Wong. “Integer programming formulations of the traveling salesman problem”. In: *Proceedings of 1980 IEEE International Conference on Circuits and Computers*. 1980, pp. 149–152.
 - [100] M. Yannakakis. “Expressing combinatorial optimization problems by linear programs”. In: *Journal of Computer and System Sciences* 43 (1991), pp. 441–466.
 - [101] G. M. Ziegler. *Lectures on polytopes*. Vol. 152. Springer Science & Business Media, 1995.

Manuel Francesco Aprile

Contact information

Office: MA B1 533, SB MATH AA Email: manuel.aprile@epfl.ch
Station 8 CH-1015 Tel.: +41 76 448 0923
EPFL, Lausanne Web: <http://disopt.epfl.ch/Aprile>
ORCID: 0000-0002-6805-6903

Research interests

Polyhedral combinatorics, combinatorial optimization, graphs and matroids

Education

From 2014: Ph.D. student in Discrete Optimization
EPFL, Lausanne, Switzerland
Advisor: Prof. Friedrich Eisenbrand
Co-advisor: Prof. Yuri Faenza (Columbia University, USA)

2014: M. Sc. in Mathematics and Foundation of Computer Science
University of Oxford, United Kingdom
Advisor: Prof. Colin McDiarmid

2013: B. Sc. in Mathematics, *summa cum laude*
University of Catania, Italy
Advisor: Prof. Giuseppe Nicosia

Awards and scholarship

2010-13: Winner of the INdAM (National Institute of High Mathematics) scholarship to entirely support my Bachelors studies (12.000 €)

Math Olympics:
2010: Bronze Medal at National Contest, Cesenatico, Italy.
2005-09: Selected for the Regional Contest, Sicily, Italy.
2004: Selected for the National Contest, Bocconi University, Milan, Italy.

Physics Olympics:
2007-09: Selected for the Regional Contest, Sicily, Italy.

Teaching

Teaching assistant at EPFL:
Spring 2015-2018: Discrete Optimization
Fall 2017: Combinatorial Geometry
Fall 2016: Combinatorial optimization
Fall 2015: Algèbre linéaire (in French)
Activities on teaching:
June 2017: SOTL workshop, Zurich (invited speaker)
Fall 2016: Science and Engineering Teaching and Learning (semester course)
May 2015: Instructional Skills Workshop (3 days)
October 2014: Teaching toolkit for Doctoral Assistants (1 day)

Student supervision	Cslovjecsek Jana:	Master project “ <i>Extension complexity of polytopes</i> ” EPFL, Fall 2018
	Gilbert Maystre:	Master project “ <i>Non-repetitive coloring of line graphs</i> ” EPFL, Spring 2017
	Loris Di Natale:	Bachelor project “ <i>On the minimum rainbow subgraph problem</i> ” EPFL, Fall 2016 (jointly with A. Cevallos)
	Maurice Amendt:	Bachelor project “ <i>Non-repetitive graph coloring</i> ” EPFL, Fall 2015
Publications	In conferences with published, peer-reviewed proceedings:	
	<ul style="list-style-type: none"> • M. Aprile, A. Cevallos, Y. Faenza (2016) <i>On Vertices and Facets of Combinatorial 2-Level Polytopes</i>. In: R. Cerulli, S. Fujishige, A. Mahjoub (eds) Combinatorial Optimization. ISCO 2016. Lecture Notes in Computer Science, vol 9849. Springer, Cham. • M. Aprile, N. Castro, F. Robledo, P. Romero (2017) <i>Analysis of Node-Resilience Strategies under Natural Disasters</i>. International Conference on Design of Reliable Communication Networks. • M. Aprile, Y. Faenza, S. Fiorini, T. Huynh, M. Macchia (2017) <i>Extension complexity of stable set polytopes of bipartite graphs</i> International Workshop on Graph-Theoretic Concepts in Computer Science. Springer, Cham, 2017. 	
	In journals:	
	<ul style="list-style-type: none"> • M. Aprile, A. Cevallos, Y. Faenza (2018) <i>On 2-level polytopes arising in combinatorial settings</i> SIAM Journal on Discrete Mathematics 32.3 • M. Aprile, N. Castro, G. Ferreira, J. Piccini, F. Robledo, P. Romero (2018) <i>Graph fragmentation problem: analysis and synthesis</i> International Transactions in Operational Research. 	
Languages	Theses:	
	<ul style="list-style-type: none"> • M. Aprile (2014) <i>Constructive Aspects of Lovász Local Lemma and Applications to Graph Colouring</i> MSc thesis, University of Oxford • M. Aprile (2013) <i>An algorithm for constructing magic squares</i> Bachelor thesis, University of Catania. 	
	Italian (native), English (proficient), French (intermediate).	