Bipartite Perfect Matching as a Real Polynomial

Gal Beniamini

Under the supervision of Professor Noam Nisan

THE HEBREW
UNIVERSITY
OF JERUSALEM

Faculty of Computer Science and Engineering
The Hebrew University of Jerusalem

A dissertation submitted to the Hebrew University of Jerusalem in partial fulfillment of the requirements for the degree of Master of Science in the Faculty of Computer Science and Engineering.

February 2020

Contents

List of Figures

Dedication

I would like to thank my advisor, Professor Noam Nisan. Thank you for introducing me to the wonderful world of Theoretical Computer Science, and for allowing me to explore it to my heart's content. I would also like to thank my parents, my brothers, and my beloved dog. Thank you for all the support, patience, and walks around the park (respectively).

Yours, Gal.

Acknowledgments

I would like to thank Professor Nati Linial for helpful discussions.

Chapter 1

Abstract

The bipartite perfect matching problem is the cornerstone of many algorithmic and complexity-theoretic tasks. Given a balanced bipartite graph, the goal is to find a collection of edges in which each vertex appears exactly once, or determine that no such set exists.

The asymptotically fastest algorithm solving the problem was found by Hopcroft and Karp [HK73], nearly half a century ago. However, despite its centrality, the complexity of bipartite perfect matching *remains undetermined*. In the regime of dense graphs, i.e., when the number of edges is $\Theta(n^2)$, the known algorithms solve the problem in $\mathcal{O}(n^{2.5})$ time. Conversely, only the "trivial" $\Omega(n^2)$ lower bound (i.e., reading the input) is known, which follows from the "evasiveness" of the problem, as shown by Yao [Yao88].

Our goal in this work is to shed more light on the complexity of bipartite perfect matching. To this end, we represent its decision variant using multilinear polynomials over the Reals. We find the explicit closed form of the unique multilinear polynomial representing the decision problem, by leveraging a connection between the polynomial and the Möbius function of the lattice of all "matching-covered" graphs. A key component of the proof is the fact that this lattice was shown by Billera and Sarangarajan [BS94] to be isomorphic to the face lattice of the Birkhoff Polytope.

In addition to the aforementioned polynomial, we also provide a fine grained characterization of its "dual" polynomial, i.e., the one in which the symbols 0 and 1 switch roles. Our proof relies heavily on the properties of the Eulerian matching-covered lattice. Crucially, we find that the number of monomials appearing in the dual polynomial is *only* exponential in $n \log n$. We consider the low number of monomials as some form of positive algorithmic result for the bipartite perfect matching problem. In particular, it implies that for the associated communication problem, the rank of the communication matrix is bounded by $2^{\mathcal{O}(n \log n)}$.

These polynomials also allow us to obtain *new lower bounds* on the problem. We extend the evasiveness result from classical decision trees, to trees whose internal nodes are labeled by XOR functions. Furthermore, we show new lower bounds for two additional families of decision trees; those whose internal nodes are labeled by AND and OR functions, respectively.

CHAPTER 2

תקציר

בעיית הזיווג המושלם הינה אחת מאבני הבניין של ענף האלגוריתמים והסיבוכיות. בהנתן גרף דו-צדדי מאוזן, נרצה לאתר קבוצת קשתות בה כל אחד מקדקדי הגרף מופיע פעם אחת בדיוק, או להכריז כי לא קיימת קבוצה כזו.

אסימפטוטית, האלגוריתם המהיר ביותר עבור בעיית הזיווג המושלם הומצא על ידי Hopcroft ו-Karp. מאז חלפו כמעט 50 שנה, ואף על פי כן, סיבוכיותה של הבעיה נותרה אפופה במסתורין. מחד גיסא, עבור גרפים בהם מספר הקשתות הוא $\Theta(n^2)$, ידועים אלגוריתמים לפתרון הבעיה בזמן $O(n^{2.5})$. מאידך גיסא, החסם התחתון היחיד הידוע הוא הטריוויאלי, $\Omega(n^2)$ (כלומר יש לקרוא את כל הקלט, שכן הבעיה "חמקנית", כפי שהראה Yao).

מטרת העל בעבודתינו היא לזרוע אור על סיבוכיותה של בעיית הזיווג המושלם. לשם כך, אנו מייצגים את בעיית ההכרעה באמצעות פולינומים מולטילינארים מעל הממשיים. בפרט, אנו מספקים ביטוי סגור עבור הפולינום הייחודי המייצג את הבעיה, תוך שימוש בקשר ההדוק שבין פונקציית מביוס של סריג הגרפים מכוסי-הזיווגים, לבין מקדמיו וגורמיו של הפולינום. הרכיב המרכזי בהוכחתינו הוא השקילות המבנית בין הסריג הנדון, לבין סריג הפאות של הפאון של Birkhoff, אשר הוכחה על ידי Billera ו-Sarangarajan.

מלבד הפולינום המתואר לעיל, אנו מספקים תיאור מפורט (אך לא מלא) של הפולינום הדואלי של בעיית הזיווג המושלם — כלומר, הפולינום בו אנו מחליפים בין תפקידם של המספרים 0 ו-1. הוכחתינו זו מסתמכת על מבנהו ותכונותיו של סריג הגרפים מכוסי-הזיווגים. תיאורו של הפולינום הדואלי מאפשר לנו להסיק כי מספר המונומים המופיעים בו הוא *רק* מעריכי ב- $n \log n$. אנו מחשיבים את מספר המונומים הנמוך בתור תוצאה אלגוריתמית חיובית חדשה עבור בעיית הזיווג המושלם. למשל, עבור בעיית התקשורת המתאימה, מספר המונומים מאפשר לנו להסיק כי דרגת מטריצת התקשורת היא לכל היותר $2^{O(n \log n)}$.

בנוסף לתוצאות האלגוריתמיות החיוביות, אנו מוכיחים בעבודתינו חסמים תחתונים חדשים לבעיית הזיווג המושלם. ראשית, אנו מראים כי תכונת ה"חמקנות" של הבעיה מתקיימת גם עבור עצי החלטה בהם כל שאילתה מתאימה לפונקציית XOR. כמו כן, אנו מראים חסמים תחתונים עבור שתי משפחות נוספות של עצי החלטה; אלו בהם השאילתות מתאימות לפונקציות AND ו-OR, בהתאמה.

Chapter 3

Introduction

Every Boolean function $f : \{0,1\}^n \to \{0,1\}$ can be represented in a unique way as a Real multilinear polynomial. This representation and related ones (e.g. using the $\{1,-1\}$ basis rather than $\{0,1\}$ – the "Fourier transform" over the hypercube, or approximation variants) have many applications for various complexity and algorithmic purposes. See, e.g., [O'D14] for a recent textbook.

In this paper we derive the representation of the bipartite-perfect-matching decision problem as a Real polynomial.

> **Definition.** *The Boolean function* $BPM_n(x_{1,1},\ldots,x_{n,n})$ *is defined to be* 1 *if and only if the bipartite graph whose edges are* $\{(i,j) : x_{i,j} = 1\}$ *has a perfect matching, and* 0 *otherwise.*

Our first result is determining the representation of this function as a Real multilinear polynomial. By way of example, $BPM_2(\bar{x}) = x_{1,1}x_{2,2} + x_{1,2}x_{2,1} - x_{1,1}x_{1,2}x_{2,1}x_{2,2}$. Somewhat surprisingly, finding the closed form expression for any n appears nontrivial. In fact, we do not know of an easier proof than our own involved proof, even showing that for any n the degree of this polynomial is n^2. [1]

To present our first result, let us introduce some notation. We will call a graph *matching-covered* if its edges can be represented as a union of perfect matchings. As an example, for $n = 2$ the graph whose edges are $\{(1,1),(1,2),(2,2)\}$ is *not* matching-covered since any perfect matching that contains the edge $(1,2)$ must also contain the edge $(2,1)$, which is not in the graph. The connected components of matching-covered graphs are called "elementary graphs" and were studied at length by [PL86]. Finally for a graph G, we denote its cyclomatic number by $\chi(G) = |E(G)| - |V(G)| + |C(G)|$ where $|C(G)|$ is the number of connected components of G. The following Theorem characterizes the multilinear polynomial of BPM_n.

[1] For the special case where n is a prime power, the full degree of the polynomial follows from an extension of the evasiveness result of [RV75], due to [NSS08]. However, for n that is not a prime power, it is not true that every monotone bipartite graph property has a full degree.

Theorem 1: The Bipartite Perfect Matching Polynomial

$$BPM_n(x_{1,1},\ldots,x_{n,n}) = \sum_{G\subseteq K_{n,n}} a_G \prod_{(i,j)\in E(G)} x_{i,j} \text{ , where:}$$

$$a_G = \begin{cases} 0 & \text{if } G \text{ is not matching-covered} \\ (-1)^{\chi(G)} & \text{if } G \text{ is matching-covered} \end{cases}$$

Our proof proceeds by studying the structure of the lattice of matching-covered graphs and its Möbius function and the key step requires using the topological structure of this lattice. Specifically, [BS94] showed that this lattice is isomorphic to the face lattice of the Birkhoff Polytope, and is thus Eulerian. Counting the number of matching-covered graphs, we get:

Corollary. *The polynomial BPM_n has $(1-o_n(1))\cdot 2^{n^2}$ monomials with non-zero coefficients.*

Our characterization of the polynomial has several corollaries. For example, it allows us to obtain a closed form expression counting the number of bipartite graphs containing a perfect matching, and in particular to show that this number is odd. It also suffices for showing that a $(1-o_n(1))$-fraction of the Fourier coefficients of BPM_n are very small, 2^{-n^2+1}, yet non-zero.

In the second part of the paper, we turn our attention towards the "dual representation" – a form in which the symbols 1 and 0 switch roles. Formally, for a Boolean function $f(x_1,\ldots,x_n)$ we define its dual by $f^\star(x_1,\ldots,x_n) = 1-f(1-x_1,\ldots,1-x_n)$. Under this notation, $BPM_n^\star(x_{1,1},\ldots,x_{n,n})$ gets the value 1 if the input graph contains a biclique over a total of $n+1$ vertices (i.e., its complement contains a violation of Hall's condition).

To present our result, we will focus on the following two classes of graphs. A bipartite graph is called *totally ordered* if there exists an ordering $v_1,\ldots,v_n$ of its left vertices such that $N(v_1)\supseteq N(v_2)\supseteq\cdots\supseteq N(v_n)$ where $N(v)$ denotes the set of right vertices connected to v. In the same vein, we call the graph *strictly totally ordered* if in fact $N(v_1)\supsetneq N(v_2)\supsetneq\cdots\supsetneq N(v_n)\supsetneq\varnothing$. For the dual case, we do not obtain a complete characterization of the polynomial. Nevertheless, we show the following fine grained characterization.

Theorem 2: The Dual Polynomial of Bipartite Perfect Matching

$$BPM_n^\star(x_{1,1},\ldots,x_{n,n}) = \sum_{G\subseteq K_{n,n}} a_G^\star \prod_{(i,j)\in E(G)} x_{i,j} \text{ , where:}$$

- If G *is not totally ordered*, we have $a_G^\star = 0$.
- If G *is strictly totally ordered*, we have $a_G^\star = (-1)^{n+1}$

Our proof relies on properties of the lattice of matching-covered graphs, and heavily utilizes its Eulerian structure. For graphs G that are totally ordered but not strictly so, the situation is complex. We show that for some such graphs G, we have $a^\star_G = 0$, for others $a^\star_G = \pm 1$, and for others still $a^\star_G \notin \{-1, 0, 1\}$. For example, for $n > 2$ and $G = K_{n-1,n-1}$ we have $a^\star_G = (n-2)^2$. We present the full polynomial of $BPM^\star_3$ in Appendix B. We leave the full characterization of the dual polynomial as an open problem.

This characterization of the dual polynomial suffices for obtaining an accurate estimate of the number of monomials with non-zero coefficients:

Corollary. *The polynomial $BPM^\star_n$ has $2^{(2n\log_2 n)\pm\mathcal{O}(n)}$ monomials with non-zero coefficients.*

We view the small number of non-zero coefficients as some form of a positive algorithmic result regarding the perfect matching problem. For example, consider a communication setting where the edges of a bipartite graph are partitioned somehow between two parties; Alice and Bob. Their task is to devise a communication protocol for determining whether the combined graph has a perfect matching. The known algorithms for bipartite matching imply a protocol that uses $\mathcal{O}(n^{1.5})$ bits of communication [DNO19, Nis19]. However, the small number of monomials in $BPM^\star_n$ directly implies that the associated communication matrix has Real rank that is *only* exponential in $n\log n$ (recall that the logarithm of the rank is a lower bound for the deterministic communication complexity, and is conjectured to be polynomially related to it).

Conversely, the polynomial representations of BPM_n and $BPM^\star_n$ allow us to obtain *new lower bounds* on the decision problem of bipartite perfect matching. Yao [Yao88] first showed that the decision problem of bipartite perfect matching is "evasive", i.e., any classical decision tree deciding it has depth exactly n^2. We extend Yao's evasiveness result to parity decision trees, wherein each internal node is labeled by an XOR function over an arbitrary subset of the input bits. In the same vein, we obtain new lower bounds for AND and OR decision trees, which are similarly defined. For each of these three families, we denote the minimal depth of a tree in the family computing BPM_n by $D^{XOR}(BPM_n)$, $D^{AND}(BPM_n)$ and $D^{OR}(BPM_n)$, respectively. We obtain the following lower bounds:

Corollary. *BPM_n is evasive for XOR decision trees, i.e., $D^{XOR}(BPM_n) = n^2$.*
Furthermore, for AND and OR decision trees, we have:

$$D^{AND}(BPM_n) \geq (\log_3 2)\cdot n^2 - o_n(1), \qquad D^{OR}(BPM_n) \geq 2\log_3(n!)$$

Of particular interest is the family of OR decision trees, which were shown by [Nis19] to be complexity preserving proxies for bipartite perfect matching; $\tilde{\mathcal{O}}(n^{1.5})$ OR queries suffice (even when slightly restricting each query), and any $\Omega(n^{1+\alpha})$ lower bound would rule out asymptotically fast algorithms from a wide class, i.e., "combinatorial algorithms".

Chapter 4

Preliminaries and Notation

4.1 Polynomial Representations of Boolean Functions

Recall the following fact regarding polynomial representations of Boolean functions (see [O'D14]):

Fact 4.1.1. *Any Boolean function $f : \{0,1\}^n \to \{0,1\}$ can be uniquely represented by a multilinear polynomial over the Reals.*

For a given a multilinear polynomial, we denote the set of all monomials appearing in it by:

Notation 4.1.2. *Let $f(x_1,\ldots,x_n) = \sum_{S\subseteq[n]} a_S \left(\prod_{i\in S} x_i\right) \in \mathbb{R}[x_1,\ldots,x_n]$ be a multilinear polynomial over the Reals. Denote the set of monomials appearing in f by:*

$$mon(f) = \{S \subseteq [n] : a_S \neq 0\}$$

4.2 The Möbius Function of Partially Ordered Sets

When discussing partially ordered sets (hereafter, **posets**), we use the Möbius function for posets. The Möbius function of a poset is the inverse, with respect to convolution, to the poset's zeta function $\zeta(y,x) = \mathbb{1}\{y < x\}$ (see, e.g., [Sta11]).

Definition 4.2.1 (Möbius Function for Posets)**.** *Let $\mathcal{P} = (P,<)$ be a finite poset. The Möbius function of the poset $\mathcal{P}$ is denoted by $\mu_P : P \times P \to \mathbb{R}$, and is defined as follows:*

$$\forall x \in P: \ \mu_P(x,x) = 1$$
$$\forall x,y \in P,\ y < x: \ \mu_P(y,x) = -\sum_{y\leq z<x} \mu_P(y,z)$$

Given a poset $\mathscr{P}$ with a unique bottom element $\hat{0}$, the values $\mu_P(\hat{0}, x)$, where $x \in \mathscr{P}$, are known as the **Möbius Numbers** of $\mathscr{P}$.

4.3 Graphs

We use the standard definitions and notation relating to graphs. For a graph G, we denote the sets of vertices and edges of G by $V(G)$ and $E(G)$, respectively. The set of all perfect matchings of G is denoted by $PM(G)$, and the set of all connected components is denoted by $C(G)$. Furthermore, for any vertex $v \in V(G)$, we denote its neighbour set by $N_G(v)$.

In addition to the quantities relating to a given graph, it will be useful to also provide some notation for basic operations on graphs. For example, the notations $G \cup \{(a,b)\}$ and $G \setminus \{(a,b)\}$ refer to the graph G with the addition or removal of the edge (a,b), respectively. In the same vein, $G - a$ is the graph where the vertex a is omitted, along with all the edges adjacent to it. Lastly, if H and G are two graphs, the notation $H \subseteq G$ indicates that $E(H) \subseteq E(G)$ and $V(H) = V(G)$.

A somewhat less common quantity which we refer to throughout the paper is the **Cyclomatic Number** of the graph, which is defined as follows:

Definition 4.3.1. *Let G be a graph. The cyclomatic number of G, $\chi(G)$, is defined:*

$$\chi(G) = |E(G)| - |V(G)| + |C(G)|$$

We will often consider the edge sets corresponding to unions of graphs. Consequently, the following notation will be useful:

Notation 4.3.2. *Let S be a set of graphs. The set of all edges appearing in any graph $G \in S$ is denoted by:*

$$\bar{E}(S) = \bigcup_{G \in S} E(G)$$

Lastly, when dealing with Boolean graph functions (i.e., Boolean functions whose input bits correspond to the edges of graphs over a fixed set of vertices), we use the following notation:

Notation 4.3.3. *Let $n, m \in \mathbb{N}^+$. Let $f : \{0,1\}^{nm} \to \{0,1\}$ be a Boolean function whose inputs are bipartite graphs over the vertices of $K_{n,m}$. Then, $\forall G \subseteq K_{n,m}$ denote:*

$$f(G) \coloneqq f(x_G), \text{ where } \forall i \in [n],\ \forall j \in [m] :\ (x_G)_{i,j} = \mathbb{1}\{(i,j) \in E(G)\}$$

4.4 Decision Trees and Query Complexity

Decision trees are binary trees whose internal nodes are labeled by Boolean functions, and whose leaves are labeled by the values $\{0,1\}$. Formally, we say that a decision tree T **computes** a Boolean function $f : \{0,1\}^n \to \{0,1\}$ if for any root-to-leaf path in T, the value of the leaf "agrees" with $f(z)$ on all inputs $z \in \{0,1\}^n$ which are *processed* by the path. An input $z \in \{0,1\}^n$ is *processed* by a path if for all functions h in the internal nodes along the path, we have $h(z) = 1$ if the path turns right at that node, and $h(z) = 0$ otherwise.

From an algorithmic perspective, decision trees can be viewed as algorithms whose every step consists of *querying* the output of some Boolean function $h \in \mathcal{H}$ on the input bits, and repeating the process until sufficient *information* is available to deduce the output. Thus, decision trees give rise to the *query complexity model*. In this model, we disregard the amount of computation required, and instead measure the minimal amount of *information*. There are several families of decision trees, which differ from one another in the set of functions $\mathcal{H}$ which label their internal nodes.

Classical Decision Trees. The most commonly studied decision trees, also known as "classical" decision trees, are those whose internal nodes are labeled by dictatorship functions, i.e., each internal node "queries" the value of a single input bit. For such trees, we use the following worst-case complexity measure:

> **Notation 4.4.1.** *Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function. The minimal depth of a classical decision tree computing f is known as the* ***Query Complexity of f****.*

Generalized Decision Trees. Three natural extensions of classical decision trees are those whose internal nodes are labeled by XOR, OR and AND functions, respectively, over arbitrary subsets of the input bits. XOR decision trees have been studied at length, and are known to be related to the Fourier expansion of a function (see, e.g., [O'D14]). OR and AND decision trees have also been studied, for example in [ML19], and in the setting of group property testing. For these three families of trees, we denote their associated *query complexities* as follows:

> **Notation 4.4.2.** *Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function. Denote the* ***minimal depth*** *of any XOR, OR or AND decision tree computing f, by $D^{XOR}(f)$, $D^{OR}(f)$ and $D^{AND}(f)$, respectively.*

4.5 Fourier Analysis

Fourier Analysis of Boolean functions is a wide field of study, in which powerful analysis tools are applied to functions over the Hamming cube, yielding combinatorial (and other) insights. Given a Boolean function $f : \{-1,1\}^n \to \{-1,1\}$, the Fourier expansion of f is the unique multilinear polynomial representing f over the Reals in the $\{1,-1\}$ basis (i.e., -1 corresponds to $True$ and 1 to $False$). The Fourier expansion of f is given by:

$$f(x_1,\ldots,x_n) = \sum_{S\subseteq[n]} \hat{f}_S \cdot \prod_{i\in S} x_i$$

Where each $\hat{f}_S$ is a Real number, referred to as the Fourier coefficient of S, and each monomial $\prod_{i\in S} x_i$ corresponds to a parity function over the set S. The aforementioned representation is unique, and the set of Fourier coefficients of f is commonly referred to as its *Fourier Spectrum*. Crucially, the set of all monomials forms an orthonormal basis. There are many important properties of the Fourier expansion, which we will not recount here. For an extensive treatment of the topic, we refer the reader to [O'D14].

Chapter 5

The Bipartite Perfect Matching Polynomial

This chapter centers around the proof of Theorem 1. We begin with some basic observations regarding a family of Boolean graph functions called "Graph Cover functions". These observations lead us to the connection between the multilinear polynomial representing BPM_n, and the Möbius numbers of the lattice of matching-covered graphs. To compute these Möbius numbers, we rely on a result of Billera and Sarangarajan [BS94], showing that the aforementioned lattice is isomorphic to the face lattice of the Birkhoff Polytope.

Using Theorem 1, we deduce several corollaries. For example, we find a closed form expression counting the number of bipartite graphs having a bipartite perfect matching, and deduce that this number is odd. We also compute asymptotically almost all the Fourier spectrum of BPM_n. Lastly, we obtain new lower bounds for decision trees; we show that BPM_n is "evasive" for XOR decision trees (i.e., exactly n^2 queries are required), and that for AND decision trees, at least $(\log_3 2) \cdot n^2 - o_n(1)$ queries are required.

5.1 Graph Cover Functions

Let $\mathcal{H}$ be a set of labeled graphs over a fixed common vertex set. Consider the following natural Boolean graph function: "Given a labeled graph G over the same vertex set, does G contain any graph in $\mathcal{H}$ as a subgraph?". In what follows, we restrict our discussion to bipartite graphs and fix our vertex set to be the vertices of the complete bipartite graph, $K_{n,m}$. Nevertheless, the same observations apply to general graphs. Formally, we define the Graph Cover function of $\mathcal{H}$ as follows:

Definition 5.1.1. *Let $\mathcal{H}$ be a set of bipartite graphs over the vertices of $K_{n,m}$. The **Graph Cover** function of $\mathcal{H}$, $f_{\mathcal{H}} : \{0,1\}^{nm} \to \{0,1\}$, is defined as follows:*

$$\forall G \subseteq K_{n,m} : \; f_{\mathcal{H}}(G) = \mathbb{1}\{\exists H \in \mathcal{H}, \; H \subseteq G\}$$

Given a set of graphs $\mathcal{H}$ over the vertices of $K_{n,m}$ and a graph $G \subseteq K_{n,m}$, we say that G is **$\mathcal{H}$-covered** if there exists some $\varnothing \neq S \subseteq \mathcal{H}$ such that $\bar{E}(S) = E(G)$. Moreover, we denote by $\mathcal{C}(\mathcal{H})$ the set of *all* $\mathcal{H}$-covered graphs. The following simple observation regarding the *monomials* of the multilinear polynomial representing $f_{\mathcal{H}}$ can be made.

Proposition 5.1.2. *Let $\mathcal{H}$ be a set of bipartite graphs over the vertices of $K_{n,m}$. The only monomials appearing in the multilinear polynomial representing $f_{\mathcal{H}}$ over the Reals are those corresponding to $\mathcal{H}$-covered graphs.*

Proof. The DNF formula representing the graph cover function is:

$$\varphi = \bigvee_{H \in \mathcal{H}} \bigwedge_{(i,j) \in E(H)} x_{i,j}$$

Since each $x_{i,j} \in \{0,1\}$, we have $\forall k \in \mathbb{N}^+$, $x_{i,j}^k = x_{i,j}$. Therefore, arithmetizing the formula yields the following polynomial representation:

$$\begin{aligned} f_{\mathcal{H}}(x_{1,1}, \ldots, x_{n,m}) &= 1 - \prod_{H \in \mathcal{H}} (1 - \prod_{(i,j) \in E(H)} x_{i,j}) \\ &= \sum_{\varnothing \neq S \subseteq \mathcal{H}} (-1)^{|S|+1} \prod_{G \in S} \prod_{(i,j) \in E(G)} x_{i,j} \\ &= \sum_{G \in \mathcal{C}(\mathcal{H})} \left(\sum_{\substack{\varnothing \neq S \subseteq \mathcal{H} \\ \bar{E}(S) = E(G)}} (-1)^{|S|+1} \right) \prod_{(i,j) \in E(G)} x_{i,j} \end{aligned}$$

□

The set of $\mathcal{H}$-covered graphs, together with the subset relation over edges, form a partially ordered set. This partially ordered set has two important properties. Firstly, it is a lattice; every two elements have a unique supremum ("join") and a unique infimum ("meet"). Secondly, the Möbius numbers of this lattice exactly describe the *coefficients* of the multilinear polynomial representing the graph cover function, $f_{\mathcal{H}}$.

Proposition 5.1.3. *Let $\mathcal{H}$ be a set of bipartite graphs over a fixed vertex set. The poset $\mathcal{P} = (\mathcal{C}(\mathcal{H}) \cup \{\hat{0}\}, \subseteq)$ is a bounded lattice, where $\hat{0}$ is the empty graph.*

Proof. The subset relation over the edges is reflexive, transitive and anti-symmetric, thus $\mathcal{P}$ is a poset. Furthermore, $\mathcal{P}$ is bounded, since $\hat{0} = (V(\mathcal{H}), \varnothing)$ and $\hat{1} = (V(\mathcal{H}), \bar{E}(\mathcal{H}))$. It remains to show that $\forall G_1, G_2 \in \mathcal{C}(\mathcal{H})$ there exists a *join* (unique supremum) and a *meet* (unique infimum).

Let $G_1, G_2 \in \mathcal{C}(\mathcal{H})$. The meet and join of G_1 and G_2 are given by:

$$E(G_1 \vee G_2) = \bigcup_{\substack{H \in \mathcal{H} \\ (H \subseteq G_1) \vee (H \subseteq G_2)}} E(H) = E(G_1) \cup E(G_2)$$
$$E(G_1 \wedge G_2) = \bigcup_{\substack{H \in \mathcal{H} \\ (H \subseteq G_1) \wedge (H \subseteq G_2)}} E(H)$$

For the join operator, let $G \coloneqq G_1 \vee G_2$. By construction, $G_1 \subseteq G$ and $G_2 \subseteq G$, therefore G is a supremum. Assume towards a contradiction that there exists another supremum $\hat{G} \neq G$ such that $G \nsubseteq \hat{G}$. Let $x \in E(G) \setminus E(\hat{G})$. Without loss of generality, assume $x \in E(G_1)$. Then $x \in E(G_1)$ and $x \notin E(\hat{G})$ therefore $G_1 \nsubseteq \hat{G}$, in contradiction to the fact that $\hat{G}$ is a supremum.

For the meet operator, let $G \coloneqq G_1 \wedge G_2$. By construction, $G \subseteq G_1$ and $G \subseteq G_2$, therefore G is an infimum. Assume towards a contradiction that there exists another infimum $\hat{G} \neq G$ such that $G \nsupseteq \hat{G}$. Let $x \in E(\hat{G}) \setminus E(G)$. Since $\hat{G} \in \mathcal{C}(\mathcal{H})$, there exists $H_x \in \mathcal{H}$ such that $H_x \subseteq \hat{G}$, $x \in E(H_x)$. However, $\hat{G}$ is an infimum, thus $H_x \subseteq \hat{G} \subseteq G_1$ and $H_x \subseteq \hat{G} \subseteq G_2$, thus by construction $H_x \subseteq G$ and $x \in E(G)$, a contradiction. □

Proposition 5.1.4. *Let $\mathcal{H}$ be a set of bipartite graphs over the vertices of $K_{n,m}$ and let $\mathcal{P} = (\mathcal{C}(\mathcal{H}) \cup \{\hat{0}\}, \subseteq)$ be the graph cover lattice of $\mathcal{H}$. Then:*

$$f_{\mathcal{H}}(x_{1,1}, \ldots, x_{n,m}) = \sum_{G \in \mathcal{C}(\mathcal{H})} -\mu_P(\hat{0}, G) \cdot \prod_{(i,j) \in E(G)} x_{i,j}$$

Namely, the coefficients of the multilinear polynomial representing $f_{\mathcal{H}}$ over the Reals are given by the (negated) Möbius numbers of $\mathcal{P}$.

Proof. Let f be the polynomial $f(x_{1,1}, \ldots, x_{n,m}) = \sum_{G \in \mathcal{C}(\mathcal{H})} -\mu_P(\hat{0}, G) \cdot \prod_{(i,j) \in E(G)} x_{i,j}$, and let $H \subseteq K_{n,m}$ be a graph. Denote by H' the union of all graphs $G \in \mathcal{C}(\mathcal{H})$ such that $G \subseteq H$. We now show that f agrees with $f_{\mathcal{H}}$ on all inputs, and deduce the identity by the uniqueness of the representing polynomial. If $H' = \hat{0}$, then indeed $f(H) = 0$ as required. Otherwise, we have:

$$\begin{aligned} f(H) &= \sum_{G \in \mathcal{C}(\mathcal{H})} -\mu_P(\hat{0}, G) \cdot \mathbb{1}\{G \subseteq H\} \\ &= \sum_{\substack{\hat{0} \subset G \subseteq H' \\ G \in \mathcal{C}(\mathcal{H})}} -\mu_P(\hat{0}, G) \\ &= \sum_{\substack{\hat{0} \subseteq G \subseteq H' \\ G \in \mathcal{C}(\mathcal{H})}} -\mu_P(\hat{0}, G) + \mu_P(\hat{0}, \hat{0}) \end{aligned}$$

And by the definition of the Möbius function, $\mu_P(\hat{0}, \hat{0}) = 1$ and $\sum_{\substack{\hat{0} \subseteq G \subseteq H' \\ G \in \mathcal{C}(\mathcal{H})}} -\mu_P(\hat{0}, G) = 0$ □

5.2 The Bipartite Perfect Matching Polynomial

Matching-Covered and Elementary Graphs

Let us begin by recalling the definition of the Boolean Bipartite Perfect Matching function:

Definition. *The Boolean Bipartite Perfect Matching function,* $BPM_n : \{0,1\}^{n^2} \to \{0,1\}$*, is defined as follows:*

$$BPM_n(x_{1,1}, \ldots, x_{n,n}) = \begin{cases} 1 & \{(i,j) : x_{i,j} = 1\} \text{ has a Perfect Matching} \\ 0 & \text{Otherwise} \end{cases}$$

The monotone Boolean function BPM_n represents the **decision problem** of bipartite perfect matching. Given a bipartite graph $G \subseteq K_{n,n}$, the function outputs 1 if and only if G contains a bipartite perfect matching. The aforementioned function may also be cast in terms of graph cover functions. In particular, it is a graph cover function for the set $\mathcal{H} = PM(K_{n,n})$. Thus, by Proposition 5.1.2, the only monomials that may appear in its multilinear polynomial over the Reals are those corresponding to $\mathcal{H}$-covered graphs. For the particular case where $\mathcal{H} = PM(K_{n,n})$, we introduce the following definition:

Definition 5.2.1. *Let* $G \subseteq K_{n,n}$ *be a balanced bipartite graph.* G *is* ***matching-covered*** *if and only if there exists some* $S \subseteq PM(K_{n,n})$ *such that* $\bar{E}(S) = E(G)$*.*

For simplicity, we introduce some notation. The set of all matching-covered graphs $H \subseteq G$ is denoted by $\mathbf{MC(G)}$. In the same vein, the set of *all* bipartite matching-covered graphs of order $2n$ is denoted $\mathbf{MC_n} := MC(K_{n,n})$. Lovász and Plummer [PL86] previously considered a family of graphs called *elementary graphs*, which are closely related to matching-covered graphs. Elementary graphs are simply the *connected components* of matching-covered graphs. Formally:

Definition 5.2.2 ([PL86])**.** G *is* ***elementary*** $\Leftrightarrow$ G *is a connected matching-covered graph.*

We recall two key Theorems regarding elementary graphs. The first, due to Hetyei [Het64], provides several necessary and sufficient conditions for elementarity of a given graph. The second, due to Lovász and Plummer [PL86], shows that all elementary graphs admit a normal form, called the *bipartite ear decomposition.*

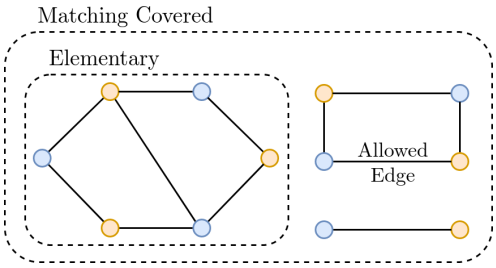

Figure 5.1: A matching-covered graph, composed of three elementary graphs

Theorem 5.2.3 ([Het64]). *Let $G = (A \cup B, E)$ be a bipartite graph. The following are equivalent:*

- *G is elementary.*
- *G has exactly two minimum vertex covers, A and B.*
- *$|A| = |B|$ and for every $\varnothing \neq X \subset A$, $|N(X)| \geq |X| + 1$.*
- *$G = K_2$, or $|V(G)| \geq 4$ and for any $a \in A$, $b \in B$, $G - a - b$ has a perfect matching.*
- *G is connected and every edge is "allowed", i.e., appears in a perfect matching of G.*

Definition 5.2.4 ([PL86]). *Let G be a balanced bipartite graph. G has a **bipartite ear decomposition** of length k if it can be written in the form:*

$$G = e + P_1 + \cdots + P_k$$

Where $e \in E(G)$, and each P_i is an odd-length path, in which any pair of adjacent vertices are from different colour classes (and in particular, so are its endpoints). The vertices appearing in each path P_i, other than its two endpoints, are "fresh" – i.e., they do not appear in $e + P_1 + \cdots + P_{i-1}$. Note that each P_i can also be a single edge connecting two preexisting vertices of different colour classes.

Theorem 5.2.5 ([PL86]). *Let G be a balanced bipartite graph. Then:*

$$G \text{ is elementary} \iff G \text{ has a bipartite ear decomposition}$$

The vast majority of balanced bipartite graphs, are in fact elementary (and in particular,

matching-covered). For example, any $G \subseteq K_{n,n}$ containing a Hamiltonian cycle is elementary – since its bipartite ear decomposition is that of the cycle, followed by single edge ears for all the remaining edges. Since asymptotically most bipartite graphs contain a Hamiltonian cycle [Fri85], the claim thus follows. Alternately, one may use Hetyei's characterization and a simple probabilistic method argument, to obtain a more explicit bound:

Proposition 5.2.6. *Let $n > 1$. Then:*

$$|MC_n| \geq |\{G \subseteq K_{n,n} : G \text{ is elementary}\}| \geq 2^{n^2}\left(1 - \frac{2n^4}{2^n}\right) = 2^{n^2}(1 - o_n(1))$$

Proof. Let $n > 1$ and let A, B be two sets, where $|A| = |B| = n$. Denote by $G(n,n,p)$ the distribution over balanced bipartite graphs of order $2n$, in which each edge appears i.i.d with probability p. Recall that by Theorem 5.2.3, $G = (A \cup B, E) \subseteq K_{n,n}$ is elementary if and only if $\forall a \in A$, $\forall b \in B$: $G - a - b$ has a perfect matching. By the union bound:

$$\begin{aligned}\Pr_{G\sim G(n,n,0.5)}[\text{G is not elementary}] &= \Pr_{G\sim G(n,n,0.5)}[\exists a \in A,\ b \in B :\ G - a - b \text{ has no perfect matching}] \\ &\leq n^2 \cdot \Pr_{G\sim G(n-1,n-1,0.5)}[\text{G has no perfect matching}]\end{aligned}$$

By Hall's Theorem, G has a perfect matching if and only if $\forall X \subseteq A$: $|N(X)| \geq |X|$. Thus G has no perfect matching if and only if there exist two sets $S \subseteq A$, $T \subseteq B$ such that $|S|+|T| = n+1$, and none of the edges in $S \times T$ appear in G. Using the union bound again:

$$\Pr_{G\sim G(n,n,0.5)}[\text{G has no perfect matching}] \leq \sum_{k=1}^{n}\binom{n}{k}\binom{n}{k-1}2^{-k(n-k+1)} \leq \frac{n^2}{2^n}$$

Thus:

$$\Pr_{G\sim G(n,n,0.5)}[G \in MC_n] \geq \Pr_{G\sim G(n,n,0.5)}[\text{G is elementary}] \geq 1 - \frac{2n^4}{2^n} \qquad \square$$

The Birkhoff Polytope and the Lattice of Matching-Covered Graphs

Let P be a polytope. The *face lattice* of P is the lattice whose elements are the faces of P, ordered by containment, together with a unique bottom element $\hat{0}$ (i.e., the "empty face") and a unique top element $\hat{1}$ (corresponding to the polytope P itself). The aforementioned lattice is *graded*, and the rank of each face $Q \neq \hat{0}$ is given by $dim(Q) + 1$.

We now recall a particular polytope: the **Birkhoff polytope**, B_n. This polytope is defined as the convex hull of all $n \times n$ permutation matrices. Billera and Sarangarajan proved the following powerful theorem regarding the face lattice of B_n:

Theorem 5.2.7 ([BS94])**.** *The face lattice of the Birkhoff polytope B_n is isomorphic to the lattice of all matching-covered graphs of order $2n$, ordered by inclusion, together with the empty graph.*

A lattice that is isomorphic to the face lattice of a polytope is known as "Eulerian". The Möbius function of an Eulerian lattice satisfies the following identity (see, e.g., [Sta11]): $\forall x \leq y : \mu(x,y) = (-1)^{rk(y)-rk(x)}$, where $rk(\cdot)$ refers to the rank of elements in the lattice. For the proof of Theorem 1, we only require the *Möbius numbers* of an Eulerian lattice. Thus, for completeness, we provide a simple proof of the identity regarding the Möbius numbers of an Eulerian lattice, using the Euler-Poincaré Formula:

Lemma 5.2.8. *Let Q be a polytope and denote by $F(Q)$ the set of all faces of Q. Let $\mathcal{P} = (F(Q) \cup \{\hat{0}\}, \leq)$ be the face lattice of Q. The Möbius numbers of $\mathcal{P}$ satisfy:*

$$\forall x \in (F(Q) \cup \{\hat{0}\}) : \ \mu_P(\hat{0}, x) = (-1)^{rk(x)}$$

Proof. Recall that every face of a polytope is also a polytope. Thus, for any face $x \in F(Q)$, we denote its face lattice by $\mathcal{P}_x$. The lattice $\mathcal{P}_x$ consists of all faces $y \in F(Q)$ where $y \leq_P x$, thus $\mathcal{P}_x$ is a sub-lattice of $\mathcal{P}$ and $\mu_{\mathcal{P}}(\hat{0}, x) = \mu_{\mathcal{P}_x}(\hat{0}, x)$. By the definition of the face lattice, the rank of any face $y \in F(x)$ in $\mathcal{P}_x$ is given by $rk(y) = dim(y) + 1$, and thus agrees with its rank in $\mathcal{P}$. Consequently, we denote the rank of any face by $rk(\cdot)$.

The proof proceeds by induction. If $x = \hat{0}$, the equality follows from the definition of the Möbius function. Otherwise, let $x \in F(Q)$, where $k := rk(x) \geq 1$. By the definition of the Möbius function and using the induction hypothesis:

$$\mu_{\mathcal{P}_x}(\hat{0}, x) = -\sum_{\substack{y \in F(x) \cup \{\hat{0}\} \\ y \neq x}} \mu_{\mathcal{P}_x}(\hat{0}, y) = -\sum_{\substack{y \in F(x) \cup \{\hat{0}\} \\ y \neq x}} (-1)^{rk(y)}$$

Since x is a Polytope of dimension $k-1$, then by the Euler-Poincaré Formula for Polytopes

(see, e.g., [Grü13]) we have:

$$\begin{aligned}1 &= \sum_{j=0}^{k-1}(-1)^j\left|\{y \in F(x) : dim(y) = j\}\right| \\ &= \sum_{x\neq y\in F(x)}(-1)^{rk(y)-1} + (-1)^{k-1} \\ &= -\sum_{\substack{y\in F(x)\cup\{\hat{0}\}\\ y\neq x}}(-1)^{rk(y)} + (-1)^{k-1} + 1 = \mu_{\mathscr{P}_x}(\hat{0}, x) + (-1)^{k-1} + 1\end{aligned}$$ □

Ranks in the Lattice of Matching-Covered Graphs

By Theorem 5.2.7, the lattice of matching-covered graphs $\mathscr{P} = (MC_n \cup \{0\}, \subseteq)$ is isomorphic to the face lattice of the Birkhoff polytope, B_n. Thus, the lattice is *graded*, and its Möbius numbers are given by $\mu_P(\hat{0}, x) = (-1)^{rk(x)}$. In a graded bounded lattice, all maximal chains have identical length. The *rank* of any element in the lattice is defined to be the length of any maximal chain in the lattice ending at the chosen element.

Using the properties of matching-covered graphs, we now show that for any matching-covered graph $G \in MC_n$, there exists a maximal chain of length $\chi(G)+1$ from G to $\hat{0}$. Therefore, we deduce that in the lattice of matching-covered graphs, we have $rk(G) = \chi(G) + 1$, where χ is the *cyclomatic number*.

Lemma 5.2.9. *Let G be an elementary graph. The following inequality holds:*

$$\forall G \neq H \in MC(G):\ \chi(H) < \chi(G)$$

Proof. Let G be an elementary graph and let $G \neq H \in MC(G)$. If H is elementary, then $|E(H)| < |E(G)|$ and $|C(H)| = |C(G)| = 1$, thus $\chi(H) < \chi(G)$, as required.

Otherwise, if H is not elementary then the connected components of H are joined by edges in G, since G is elementary and in particular connected. Observe that every connected component of H must be incident to at least 2 edges in $E(G) \setminus E(H)$ – one connected to a left vertex of the connected component, and another to a right vertex. Assume toward a contradiction that this were not the case, then there exists a component $C \in C(H)$ which is only incient to a single edge $e \in (E(G) \setminus E(H))$. Since G is elementary, there exists some perfect matching involving e (i.e., the edge e is allowed). However, upon selecting the edge e, the component C becomes *unbalanced*, and therefore the perfect matching cannot be extended over C, a contradiction.

Thus, since each component of G has at least two incident edges in $E(G) \setminus E(H)$, we have that $|E(H)| + |C(H)| \leq |E(G)|$ (i.e., if the incident edges form a cycle joining the components

$C(H)$). Thus:

$$\chi(H) = |E(H)| - |V(H)| + |C(H)| \leq |E(G)| - |V(G)| < \chi(G)$$ □

Corollary 5.2.9.1. *Let $G \in MC_n$. The following inequality holds:*

$$\forall G \neq H \in MC(G): \ \chi(H) < \chi(G)$$

Proof. If H is elementary, the proof follows from Lemma 5.2.9. Otherwise, the proof follows from the additivity of χ, by applying Lemma 5.2.9 to each connected component in which G and H differ. □

Lemma 5.2.10. *Let $G \in MC_n$, $G \notin PM(K_{n,n})$. Then there exists $H \in MC(G)$ such that:*

$$\chi(H) = \chi(G) - 1$$

Proof. Let $G \in MC_n$, $G \notin PM(K_{n,n})$. Since G is not a perfect matching, there exists a component $C \in C(G)$ such that $C \neq K_2$. C is elementary, and therefore there exists a bipartite ear decomposition: $C = e + P_1 + \cdots + P_k$. Let $C' = e + P_1 + \ldots P_{k-1}$, and observe that since C' has a bipartite ear decomposition, it too is elementary.

If P_k is a single edge, then we construct H by taking G, and replacing the component C with C'. Observe that $H \in MC_n$, and furthermore $|C(H)| = |C(G)|$ and $|E(H)| = |E(G)| - 1$. Thus $\chi(H) = \chi(G) - 1$, as required.

Otherwise, P_k is an ear $(v_1, u_1, \ldots, v_t, u_t)$. In this case, we construct H by taking G, replacing C with C', and replacing the ear P_k with the edges $(v_2, u_1), \ldots, (v_t, u_{t-1})$. Once again, $H \in MC_n$ (since all its components are elementary). Furthermore $|C(H)| = |C(G)| + t - 1$ and $|E(H)| = |E(G)| - t$, and thus $\chi(H) = \chi(G) - 1$. □

Thus, combining Corollary 5.2.9.1 and Lemma 5.2.10, we find that:

Corollary 5.2.10.1. *Let $\mathcal{P} = (MC_n \cup \{\hat{0}\}, \subseteq)$ be the lattice of matching-covered graphs. Then:*

$$\forall \hat{0} \neq G \in MC_n: \quad rk(G) = \chi(G) + 1$$

Completing the Proof of Theorem 1

We are now ready to prove the main theorem for this section:

> **Theorem 1.** *The unique multilinear polynomial representing BPM_n over the Reals is:*
>
> $$BPM_n(x_{1,1},\ldots,x_{n,n}) = \sum_{G\in MC_n} (-1)^{\chi(G)} \cdot \prod_{(i,j)\in E(G)} x_{i,j}$$

Proof. Let $\mathcal{P} = (MC_n \cup \{\hat{0}\}, \subseteq)$ be the lattice of matching-covered graphs, and let B_n be the Birkhoff Polytope. Since BPM_n is a graph cover function for the set $PM(K_{n,n})$, then by Proposition 5.1.4 we have:

$$BPM_n(x_{1,1},\ldots,x_{n,n}) = \sum_{G\in MC_n} -\mu_P(\hat{0},G) \cdot \prod_{(i,j)\in E(G)} x_{i,j}$$

By Theorem 5.2.7, $\mathcal{P}$ is isomorphic to the face lattice of B_n, and thus by Corollary 5.2.10.1 and Lemma 5.2.8, we get:

$$\forall G \in MC_n : \ \mu_P(\hat{0},G) = (-1)^{rk(G)} = (-1)^{\chi(G)+1}$$

□

5.3 Another Technique for Evasiveness?

The proof regarding the multilinear polynomial of BPM_n could, perhaps, be viewed as another "technique" for evasiveness. Given a (not necessarily bipartite) graph cover function whose corresponding lattice is isomorphic to the face lattice of some polytope, we can conclude that the function has full polynomial degree over the Reals, and is thus evasive. In fact, such functions also exhibit full polynomial degree over $\mathbb{F}_2$, and are therefore evasive even for XOR decision trees. Nevertheless, we are presently only aware of two such functions exhibiting an isomorphism between their lattice and the face lattice of a polytope – the first being that of bipartite perfect matching and the Birkhoff polytope, and the second being the OR_n function and the n-dimensional Hypercube.

Previously, Kahn, Saks and Sturtevant [KSS84] showed a topological approach for evasiveness of monotone graph properties. Given a monotone graph property $\mathcal{P}$, their technique considers the abstract simplicial complex formed by all sets in the complement of $\mathcal{P}$, and shows that if the aforementioned complex is not contractible, then the property is evasive.

These two techniques are incomparable. While the [KSS84] technique is much more widely applicable, it does not imply that monotone graph properties exhibit full polynomial degree, neither over $\mathbb{F}_2$ nor over the Reals (and indeed, many do not). Nevertheless, our approach for

evasiveness appears useful only in cases where the Möbius numbers of the corresponding lattice are "easy" to compute, e.g. when the lattice is isomorphic to the face lattice of a polytope. Therefore, this technique appears rather limited.

5.4 Corollaries of Theorem 1

Corollary 5.4.0.1. *The number of monomials in BPM_n is at least $2^{n^2}\left(1-\frac{2n^4}{2^n}\right)$.*

Proof. The bound follows immediately from Proposition 5.2.6 and Theorem 1. □

Corollary 5.4.0.2. *The degree of BPM_n over the Reals, as well as over $\mathbb{F}_2$, is n^2.*

Corollary 5.4.0.3. *BPM_n is evasive, even for XOR decision trees:*

$$D^{XOR}(BPM_n) = n^2$$

Proof. We show that for any Boolean function f, $D^{XOR}(f) \geq deg_2(f)$, where $deg_2(f)$ is the degree of the polynomial representing f over $\mathbb{F}_2$. Thus in particular $D^{XOR}(BPM_n) = n^2$. Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function and let T be a XOR decision tree computing f. Let $\mathscr{P}$ be the set of all root to 1-leaf paths in T. For any path $P \in \mathscr{P}$ we construct the indicator over the path, denoted $\mathbb{1}_P(x_1,\ldots,x_n)$, by taking the product over any parity along the path (taking the parity itself for any right turn, and adding 1 to the term for any left turn). Observe that $f(x_1,\ldots,x_n) = \sum_{P\in\mathscr{P}} \mathbb{1}_P(x_1,\ldots,x_n)$, therefore $deg_2(f) \leq \max_{P\in\mathscr{P}} deg_2(\mathbb{1}_P(x_1,\ldots,x_n)) \leq depth(T)$, where the last inequality follows since parities over $\mathbb{F}_2$ are linear functionals. □

Corollary 5.4.0.4. *The number of balanced bipartite graphs of order $2n$ containing a perfect matching is odd. Furthermore, the number of matching-covered graphs of order $2n$ is also odd.*

Proof. For any Boolean function f, $|\{x \in \{0,1\}^n : f(x) = 1\}| \equiv 1 \pmod 2$ if and only if the polynomial representing f over $\mathbb{F}_2$ has full degree. Thus the number of graphs containing a perfect matching is odd. Let $H \in MC_n$. Clearly H has a perfect matching, therefore:

$$\begin{aligned} 1 = BPM_n(H) &= \sum_{G\in MC_n} (-1)^{\chi(G)} \cdot \mathbb{1}\{G \subseteq H\} \\ &= \sum_{G\in MC(H)} (-1)^{\chi(G)} \end{aligned}$$

In particular $K_{n,n} \in MC_n$, thus:

$$1 \equiv BPM_n(K_{n,n}) \pmod 2 \equiv |MC_n| \pmod 2 \qquad \square$$

Lower Bound for AND Decision Trees

Given a Boolean function $f : \{0,1\}^n \to \{0,1\}$, the unique multilinear polynomial representing f can be used to deduce lower bounds on the query complexity of f for AND decision trees.

Lemma 5.4.1. *Let $f : \{0,1\}^n \to \{0,1\}$ be a boolean function. Then:*

$$D^{AND}(f) \geq \log_3(|mon(f)|)$$

Proof. Let T be an AND decision tree computing f and denote $d = depth(T)$. Let $\mathcal{P}$ be the set of all root to 1-leaf paths in T. For any $P \in \mathcal{P}$, construct the indicator function for the path as follows:

$$\mathbb{1}_P(x_1, \ldots, x_n) = \left(\prod_{\neg AND(S) \in P} \left(1 - \prod_{i \in S} x_i\right)\right)\left(\prod_{AND(S) \in P} \left(\prod_{i \in S} x_i\right)\right)$$

Notice that the multilinear polynomial of each indicator function of a path P making k left turns has at most 2^k monomials. Furthermore, in a binary tree of depth d there are at most $\binom{d}{k}$ paths making exactly k left turns (i.e., by selecting the position in the path at which the left turns are made). Finally, observe that $f(x_1, \ldots, x_n) = \sum_{P \in \mathcal{P}} \mathbb{1}_P(x_1, \ldots, x_n)$. Thus by the uniqueness of the multilinear polynomial representing f, we have:

$$|mon(f)| \leq \sum_{P \in \mathcal{P}} |mon(\mathbb{1}_P)| \leq \sum_{k=0}^{d} \binom{d}{k} 2^k = 3^d \qquad \square$$

Applying the aforementioned lemma to BPM_n (and recalling Corollary 5.4.0.1), we obtain:

Corollary 5.4.1.1. *The depth of any AND decision tree computing BPM_n is at least:*

$$D^{AND}(BPM_n) \geq (\log_3 2) \cdot n^2 + o_n(1)$$

We note that the same decision tree lower bound can in fact be derived for the l_1-norm of the coefficient vector of the multilinear polynomial, rather than the number of monomials. Generally, the l_1 norm provides a stronger lower bound, however for the case of BPM_n this does not yield a better bound, since all its coefficients have a magnitude of exactly 1.

5.5 The Fourier Spectrum of BPM_n

In this section we briefly discuss another multilinear polynomial representing BPM_n over the Reals – the Fourier Expansion of BPM_n. Given a multilinear polynomial over the Reals representing a Boolean function f in the $\{0,1\}$ basis, the polynomial can be "converted" into the Fourier expansion of f by replacing each monomial $\prod_{i\in S} x_i$ with the indicator $\mathbb{1}_S(x_1,\ldots,x_n) = \prod_{i\in S} \frac{-x_i+1}{2}$, and applying the transformation $x \mapsto -2x+1$ to the output. Thus:

Lemma 5.5.1. *Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function represented by the Real multilinear polynomial:*

$$f(x_1,\ldots,x_n) = \sum_{S\subseteq[n]} a_S \cdot \prod_{i\in S} x_i$$

Then the Fourier expansion of f is given by:

$$\hat{f}(x_1,\ldots,x_n) = 1 + \sum_{S\subseteq[n]} \left((-1)^{|S|-1} \sum_{T\supseteq S} \frac{a_T}{2^{|T|-1}} \right) \cdot \prod_{i\in S} x_i$$

Combining Lemma 5.5.1 and Theorem 1, we thus conclude that:

Corollary 5.5.1.1. *The Fourier coefficients of BPM_n are given by:*

$$\forall \hat{0} \neq G \subseteq K_{n,n} : \quad \widehat{BPM_n^G} = (-1)^{|E(G)|} \sum_{\substack{H\supseteq G \\ H\in MC_n}} \frac{(-1)^{\chi(H)-1}}{2^{|E(H)|-1}}$$

While the above expression might be difficult to compute in the general case, we will now see that for the asymptotic majority of graphs (all elementary graphs), the Fourier coefficient can be exactly computed.

Proposition 5.5.2. *Let $n \in \mathbb{N}^+$ and let $G \subseteq K_{n,n}$ be an elementary graph. Then:*

$$\widehat{BPM_n^G} = 2^{-n^2+1}$$

Proof. If G is elementary, then any graph $H \supseteq G$ is also elementary, as its ear decomposition is that of G, followed by adding single-edge ears for each edge in $E(H) \setminus E(G)$. Thus by Theorem

1 and Lemma 5.5.1:

$$\widehat{BPM_n^G} = (-1)^{|E(G)|-1} \sum_{\substack{H \supseteq G \\ H \in MC_n}} \frac{(-1)^{\chi(H)}}{2^{|E(H)|-1}}$$
$$= \sum_{t=0}^{n^2-|E(G)|} \binom{n^2-|E(G)|}{t} \frac{(-1)^t}{2^{t+|E(G)|-1}} = 2^{-n^2+1}$$ □

Corollary 5.5.2.1. *Let $n > 0$. For $(1 - o_n(1)) \cdot 2^{n^2}$ of the Fourier coefficients, we have:*

$$\widehat{BPM_n^G} = 2^{-n^2+1}$$

Proof. By Proposition 5.2.6, the number of elementary graphs is at least:

$$|\{G \subseteq K_{n,n} : \text{G is elementary}\}| \geq \left(1 - \frac{2n^4}{2^n}\right) \cdot 2^{n^2} = (1 - o_n(1)) \cdot 2^{n^2}$$ □

Like all monotone graph properties, bipartite perfect matching exhibits a sharp threshold. Erdős and Rényi first showed [ER64] that in the random graph model $G(n, n, p)$, when $p = \frac{f(n)+\ln n}{n}$, then the probability that a perfect matching exists satisfies:

$$\lim_{n\to\infty} \Pr_{G\sim G(n,n,p)}[G \text{ has a Perfect Matching}] = \begin{cases} 0 & f(n) \to -\infty \\ e^{-2e^{-c}} & f(n) \to c \\ 1 & f(n) \to \infty \end{cases}$$

Using Hall's Theorem, it is not hard to show that for $p = 1/2$, the probability that a perfect matching exists in a randomly sampled bipartite graph is $1 - \frac{\text{poly}(n)}{2^n}$. We now use the Fourier expansion of BPM_n in order to derive an explicit closed-form expression for the aforementioned probability, in terms of matching-covered graphs and their cyclomatic number.

Proposition 5.5.3. *Let $n > 0$. The probability that a perfect matching exists in a uniformly sampled balanced bipartite graph of order $2n$ is:*

$$\Pr_{G\sim G(n,n,1/2)}[G \text{ has a Perfect Matching}] = \sum_{G \in MC_n} \frac{(-1)^{\chi(G)}}{2^{|E(G)|}}$$

Proof. By Theorem 1 and Lemma 5.5.1, the Fourier coefficient of the empty set in BPM_n is:

$$\widehat{BPM_n^{\varnothing}} = 1 - \sum_{G \in MC_n} \frac{(-1)^{\chi(G)}}{2^{|E(G)|-1}}$$

Furthermore, for any Boolean function $f : \{1,-1\}^n \to \{1,-1\}$:

$$\hat{f}_{\varnothing} = \mathop{\mathbb{E}}_{x \sim \{1,-1\}^n}[f(x)] = \mathop{\mathbb{E}}_{x \sim \{1,-1\}^n}[-2 \cdot \mathbb{1}\{f(x) = -1\} + 1] = -2 \cdot \Pr_{x \sim \{1,-1\}^n}[f(x) = -1] + 1$$

And the equality now follows by rearranging. □

Chapter 6

The Dual Bipartite Perfect Matching Polynomial

In the previous chapter, we dealt with the unique multilinear polynomial representing BPM_n over the Reals in the $\{0,1\}$ basis (Section 5.2). We also briefly encountered the multilinear polynomial representing BPM_n in the $\{1,-1\}$ basis, i.e., its Fourier expansion (Section 5.5).

We now turn our attention towards a *third* multilinear polynomial, the one representing the "dual function" of BPM_n; namely, the function in which the symbols 0 and 1 have been "flipped", whereby 1 indicates $False$ and 0 is $True$ (which we refer to as the "$\{1,0\}$-basis"). In this chapter we will prove Theorem 2, which exhibits a fine-grained characterization of the dual polynomial. To this end, let us now introduce several more useful definitions and notation for this chapter.

6.1 Definitions and Notation

Dual Functions

Definition 6.1.1. *Let $f:\ \{0,1\}^n \to \{0,1\}$ be a Boolean function. The **dual function** of f, denoted $f^\star:\ \{0,1\}^n \to \{0,1\}$, is defined by:*

$$\forall x \in \{0,1\}^n:\ f^\star(x_1,\ldots,x_n) = 1 - f(1-x_1,\ldots,1-x_n)$$

Hereafter, we denote by $\mathbf{BPM_n^\star}$ the dual function of BPM_n. For any graph $G \subseteq K_{n,n}$, we denote its corresponding coefficient in the polynomial representing $BPM_n^\star$ by $\mathbf{a_G^\star}$. Under this notation, the polynomial representing $BPM_n^\star$ is given by:

$$BPM_n^\star(x_{1,1},\ldots,x_{n,n}) = \sum_{G\subseteq K_{n,n}} a_G^\star \cdot \prod_{(i,j)\in E(G)} x_{i,j}$$

Graphs

Hall's Theorem states that a balanced bipartite graph G has *no perfect matching* if and only if there exists an anticlique over a total of $n+1$ vertices. In this chapter, we consider dual functions, wherein the input bits (and output bit) are flipped. Thus, it will be useful to consider the "dual" to the above condition: the set of all *complete bipartite graphs* over a total of $n+1$ vertices. We will hereafter refer to any such graph as a "Hall Violator", and will use the following notation for these graphs and for graphs which are "covered" by them:

Notation 6.1.2. *Let $n > 1$. The set of all "Hall Violator" graphs is defined as follows:*

$$\mathbf{HV_n} = \{K_{X,Y} \subseteq K_{n,n} : |X| + |Y| = n + 1\}$$

Where $K_{X,Y}$ is the complete bipartite graph whose edges are $X \times Y$, and the remaining vertices are isolated.

Notation 6.1.3. *Let $n > 1$. The set of all graphs which are "covered" by Hall violators is denoted by $\mathbf{HVC_n}$, where for every $G \subseteq K_{n,n}$:*

$$G \in HVC_n \iff \exists S \subseteq HV_n : \bar{E}(S) = E(G)$$

We also consider the following two families of graphs:

Definition 6.1.4. *Let $n > 1$. A bipartite graph $G \subseteq K_{n,n}$ is called **totally ordered** if there exists an ordering of its left vertices $\{a_1, \ldots, a_n\}$, such that:*

$$N_G(a_1) \supseteq N_G(a_2) \supseteq \cdots \supseteq N_G(a_n)$$

*Similarly, G is called **strictly totally ordered** if in fact:*

$$N_G(a_1) \supsetneq N_G(a_2) \supsetneq \cdots \supsetneq N_G(a_n) \supsetneq \varnothing$$

6.2 A Fine Grained Characterization of the Dual Polynomial

In this section, we obtain a fine grained characterization of the multilinear polynomial representing $BPM_n^\star$. Unlike the multilinear polynomial of BPM_n, we do not provide an explicit

closed form of this polynomial. Nevertheless, we obtain an asymptotically tight estimate of the number of monomials appearing in the dual polynomial. Our characterization is the following:

Theorem 2. *Let $n > 1$ and let $BPM_n^\star$ be the dual function of BPM_n, represented by the following multilinear polynomial over the Reals:*

$$BPM_n^\star(x_{1,1}, \ldots, x_{n,n}) = \sum_{G \subseteq K_{n,n}} a_G^\star \prod_{(i,j) \in E(G)} x_{i,j}$$

Then for every $G \subseteq K_{n,n}$, we have:

- *If G is not totally ordered, then $a_G^\star = 0$.*
- *If G is strictly totally ordered, then $a_G^\star = (-1)^{n+1}$*

For the remainder of this section, we will set about proving Theorem 2.

$\mathbf{BPM_n^\star}$ as a Graph Cover Function

Let $G \subseteq K_{n,n}$. By Hall's Theorem, G has a perfect matching if and only if its complement does not have a biclique over $n+1$ vertices. Therefore, by the definition of the dual function, we have:

$$\begin{aligned} BPM_n^\star(G) &= \mathbb{1}\{\bar{G} \text{ does not have a perfect matching}\} \\ &= \mathbb{1}\{\bar{G} \text{ has an anticlique over a total of } n+1 \text{ vertices}\} \\ &= \mathbb{1}\{G \text{ has a biclique over a total of } n+1 \text{ vertices}\} \\ &= \mathbb{1}\{\exists H \in HV_n,\ H \subseteq G\} \end{aligned}$$

Thus, $BPM_n^\star$ is a graph cover function over the set HV_n. In particular, by Proposition 5.1.2, the only monomials appearing in the multilinear polynomial representing $BPM_n^\star$ are those corresponding to graphs $G \in HVC_n$.

This observation alone already restricts the possible graphs which may appear as monomials of $BPM_n^\star$. For example, it allows us to deduce that every $G \subseteq K_{n,n}$ with $a_G^\star \neq 0$ has a single non-trivial connected component, since every $H \in HV_n$ appearing in G contributes a connected component with exactly $n+1$ vertices. Nevertheless, this restriction does not suffice for bounding the number of monomials of $BPM_n^\star$ (as is exemplified later, in Subsection 6.3). Thus we now turn to our second characterization.

Using The Eulerian Lattice of Matching-Covered Graphs

The characterization of $BPM_n^\star$ as a graph cover function for the lattice of graphs covered by Hall violators allowed us to restrict the set of *monomials* that may appear in its polynomial representation. To gain further headway, we now shift our attention back to the Eulerian lattice of matching-covered graphs. Ideally, it would be advantageous to take the "neat" representation of BPM_n in terms of the lattice of matching-covered graphs, and "convert" it into a characterization of $BPM_n^\star$.

Given a multilinear polynomial over the Reals representing any Boolean function $f : \{0,1\}^n \to \{0,1\}$, the dual polynomial of f can immediately be derived by negating the inputs and output of f. In particular, if the polynomial representing f is given by:

$$f(x_1,\ldots,x_n) = \sum_{S\subseteq[n]} a_S \cdot \prod_{i\in S} x_i$$

Then the dual polynomial of f can be expressed as follows:

$$\begin{aligned} f^\star(x_1,\ldots,x_n) &= 1 - f(1-x_1,\ldots,1-x_n) = 1 - \sum_{T\subseteq[n]} a_T \cdot \prod_{i\in T}(1-x_i) \\ &= 1 - \sum_{T\subseteq[n]} a_T \left(\sum_{S\subseteq T} (-1)^{|S|} \prod_{i\in S} x_i \right) \\ &= 1 + \sum_{S\subseteq[n]} (-1)^{|S|+1} \left(\sum_{T\supseteq S} a_T \right) \cdot \prod_{i\in S} x_i \end{aligned}$$

Thus by applying the above to BPM_n and using the characterization of Theorem 1, we obtain the following:

Lemma 6.2.1. *Let $\mathscr{P} = (MC_n \cup \{\hat{0}\}, \subseteq)$ be the lattice of matching-covered graphs. Then for every nonempty $G \subseteq K_{n,n}$, we have:*

$$a_G^\star = (-1)^{|E(G)|+1} \sum_{\substack{G\subseteq H\subseteq K_{n,n} \\ H\in MC_n}} (-1)^{\chi(H)} = (-1)^{|E(G)|} \sum_{\substack{G\subseteq H\subseteq K_{n,n} \\ H\in MC_n}} \mu_P(\hat{0}, H)$$

We now show the following powerful characterization, which leverages the properties of the Möbius function of an Eulerian lattice:

Lemma 6.2.2. *Let $n > 1$. For all $G \in (MC_n \setminus \{K_{n,n}\})$, we have $a_G^\star = 0$.*

Proof. Let $G \in (MC_n \setminus \{K_{n,n}\})$ and let $\mathscr{P} = (MC_n \cup \{\hat{0}\}, \subseteq)$ be the Eulerian lattice of matching-covered graphs, where $\hat{0}$ is the empty graph. Since $\mathscr{P}$ is Eulerian, its Möbius function satisfies

the following identity [Sta11]: $\forall H \in MC_n, H \supseteq G:\ \mu_{\mathscr{P}}(\hat{0}, H) = \mu_{\mathscr{P}}(\hat{0}, G) \cdot \mu_{\mathscr{P}}(G, H)$. Therefore by Definition 4.2.1 and Lemma 6.2.1, we have:

$$\begin{aligned} a^{\star}_{G} &= (-1)^{|E(G)|} \sum_{\substack{G \subseteq H \subseteq K_{n,n} \\ H \in MC_n}} \mu_P(\hat{0}, H) \\ &= (-1)^{|E(G)|} \mu_{\mathscr{P}}(\hat{0}, G) \sum_{\substack{G \subseteq H \subseteq K_{n,n} \\ H \in MC_n}} \mu_{\mathscr{P}}(G, H) = 0 \end{aligned}$$ □

Extending Beyond the Lattice of Matching-Covered Graphs

By using the properties of the Möbius function of an Eulerian lattice, we were able to deduce that all matching-covered graphs (other than the complete bipartite graph), have a zero dual coefficient. While matching-covered graphs constitute the asymptotic majority of all balanced bipartite graphs, the previous observation is nevertheless insufficiently powerful to obtain our bound (indeed there are at least 2^{n^2-2n} graphs that are not matching-covered).

Subsequently, we now extend our characterization to graphs beyond the lattice of matching-covered graphs. To this end, we introduce the notion of "umbrellas" – a set of matching-covered graphs that forms a "basis" for a given graph G, even when G itself is not matching-covered.

Notation 6.2.3. *Let $n > 1$ and let $G \subseteq K_{n,n}$ be a graph. The **Umbrella of G**, $\mathcal{U}(G) \subseteq MC_n$, is the set of all minimal matching-covered graphs, with respect to containment, which contain G as a subgraph. Formally:*

$$H \in \mathcal{U}(G) \iff (G \subseteq H \in MC_n) \wedge (\nexists\, H' \in MC_n:\ G \subseteq H' \subset H)$$

The umbrella of G is an anti-chain in the lattice of matching-covered graphs. In particular, any matching-covered graph $H \in MC_n$ containing G as a subgraph, also contains a graph from the umbrella of G. Using umbrellas we now show the following identity for general graphs (i.e., not necessarily matching-covered):

Lemma 6.2.4. *Let $n > 1$ and let $G \subseteq K_{n,n}$ be a nonempty graph. Then:*

$$a^{\star}_{G} = (-1)^{n+|E(G)|} \cdot \sum_{\substack{\varnothing \neq S \subseteq \mathcal{U}(G) \\ \bar{E}(S) = K_{n,n}}} (-1)^{|S|+1}$$

Proof. Let $n > 1$ and let $\mathcal{P} = (MC_n \cup \{\hat{0}\}, \subseteq)$ be the Eulerian lattice of matching-covered graphs, where $\hat{0}$ is the empty graph. Let $G \subseteq K_{n,n}$, where $G \neq \hat{0}$. For any $\varnothing \neq S \subseteq MC_n$, denote by $\bigvee S$ the *join* of all graphs in S. Recall (Proposition 5.1.3) that in $\mathcal{P}$, the join $\bigvee S$ is the union of all graphs in S. By Lemma 6.2.1, we have:

$$a_G^\star = (-1)^{|E(G)|} \cdot \sum_{\substack{G \subseteq H \subseteq K_{n,n} \\ H \in MC_n}} \mu_{\mathcal{P}}(\hat{0}, H)$$

Using the inclusion-exclusion principle on the umbrella of G, $\mathcal{U}(G)$, we obtain:

$$\begin{aligned}\sum_{\substack{G \subseteq H \subseteq K_{n,n} \\ H \in MC_n}} \mu_{\mathcal{P}}(\hat{0}, H) &= \sum_{\varnothing \neq S \subseteq \mathcal{U}(G)} (-1)^{|S|+1} \sum_{\substack{(\bigvee S) \subseteq H \subseteq K_{n,n} \\ H \in MC_n}} \mu_{\mathcal{P}}(\hat{0}, H) \\ &= \left(\sum_{\substack{\varnothing \neq S \subseteq \mathcal{U}(G) \\ (\bigvee S) \subset K_{n,n}}} (-1)^{|S|+1} \sum_{\substack{(\bigvee S) \subseteq H \subseteq K_{n,n} \\ H \in MC_n}} \mu_{\mathcal{P}}(\hat{0}, H) + \sum_{\substack{\varnothing \neq S \subseteq \mathcal{U}(G) \\ (\bigvee S) = K_{n,n}}} (-1)^{|S|+1} \mu_{\mathcal{P}}(\hat{0}, K_{n,n}) \right)\end{aligned}$$

Since $\mathcal{P}$ is Eulerian, the sum of Möbius numbers in any **nontrivial closed interval** is zero (see Lemma 6.2.2). In particular, for any $S \subseteq \mathcal{U}(G)$ where $(\bigvee S) \neq K_{n,n}$, we have:

$$\sum_{\substack{(\bigvee S) \subseteq H \subseteq K_{n,n} \\ H \in MC_n}} \mu_{\mathcal{P}}(\hat{0}, H) = 0$$

Therefore:

$$\begin{aligned}a_G^\star &= (-1)^{|E(G)|} \cdot \sum_{\substack{\varnothing \neq S \subseteq \mathcal{U}(G) \\ \bar{E}(S) = K_{n,n}}} (-1)^{|S|+1} (-1)^{\chi(K_{n,n})+1} \\ &= (-1)^{n+|E(G)|} \cdot \sum_{\substack{\varnothing \neq S \subseteq \mathcal{U}(G) \\ \bar{E}(S) = K_{n,n}}} (-1)^{|S|+1}\end{aligned}$$

□

Given a graph $G \subseteq K_{n,n}$, we say that G has an **Incomplete Umbrella** if $\bar{E}(\mathcal{U}(G)) \neq K_{n,n}$, i.e., there exists some edge which is not present in any of the graphs in the umbrella of G. Observe that by Lemma 6.2.4, this is a sufficient condition for exhibiting a zero dual coefficient.

Corollary 6.2.4.1. *Let $n > 1$ and let $G \subseteq K_{n,n}$ be a nonempty graph. Then:*

$$\bar{E}(\mathcal{U}(G)) \neq K_{n,n} \implies a_G^\star = 0$$

Graphs with an Incomplete Umbrella

Definition 6.2.5. *Let $n > 1$ and let $G \subseteq K_{n,n}$ be a nonempty graph. An edge $(a,b) \notin E(G)$ is called a **Wildcard Edge** for G if and only if:*

$$\forall H \in MC_n,\ H \supseteq G \cup \{(a,b)\} :\ (H \setminus \{(a,b)\}) \in MC_n$$

Lemma 6.2.6. *Let $n > 1$ and let $G \subseteq K_{n,n}$ be a nonempty graph. Then:*

$$G \text{ has a wildcard edge} \implies G \text{ has an incomplete umbrella}$$

Proof. Let $(a,b) \notin E(G)$ be a wildcard edge for G. We show that $(a,b) \notin \bar{E}(\mathcal{U}(G))$. Assume towards a contradiction that $(a,b) \in \bar{E}(\mathcal{U}(G))$, and let $H \in \mathcal{U}(G)$ be a graph such that $(a,b) \in E(H)$. Then by the definition of (a,b), we have $H' = (H \setminus \{(a,b)\}) \in MC_n$, and furthermore $G \subseteq H' \subset H$, in contradiction to the fact that $H \in \mathcal{U}(G)$. □

Building upon wildcard edges, we now introduce the following (slightly weaker) sufficient condition:

Definition 6.2.7. *Let $n > 1$ and let $G \subseteq K_{n,n}$ be a nonempty graph. Denote by A the set of left vertices of G. An edge $(a,b) \notin E(G)$ is called a **Surplus Edge** for G if and only if:*

$$\forall X \subset A,\ a \in X,\ b \notin N_G(X) :\ |N_G(X)| > |X|$$

The above can be seen as a strengthening of Hall's condition, in which we require that the condition holds with a *positive surplus*. However, note that we only require the condition for a particular family of sets – those in which a is present, and b is not in the neighbour set. Finally, we show that surplus edges are, in fact, wildcard edges.

Lemma 6.2.8. *Let $n > 1$ and let $G \subseteq K_{n,n}$ be a nonempty graph. Then:*

$$(a,b) \notin E(G) \text{ is a surplus edge for } G \implies (a,b) \notin E(G) \text{ is a wildcard edge for } G$$

Proof. Let $(a,b) \notin E(G)$ be a surplus edge for G and let $H \in MC_n$ such that $H \supseteq G \cup \{(a,b)\}$. Denote $H' = H \setminus \{(a,b)\}$. It remains to show that $H' \in MC_n$. Assume towards a contradiction

that $H' \notin MC_n$ and denote by $C = (A_C \cup B_C, E) \in C(H)$ the connected component of H containing the edge (a,b). First, note that $K_2 \neq (C \setminus \{(a,b)\}) \in C(H')$, since elementary graphs are $2-$*connected*. Since $C \setminus \{(a,b)\}$ is not elementary, then by Theorem 5.2.3 there exists $\varnothing \neq X \subset A_C \subseteq A$ such that $|N_{H'}(X)| \leq |X|$.

Observe that $a \in X$ and $b \notin N_{H'}(X)$. Otherwise, we have $N_{H'}(X) = N_H(X)$ and since $H \in MC_n$ then C is elementary and thus $|N_{H'}(X)| = |N_H(X)| > |X|$, a contradiction. However, since (a,b) is a surplus edge for G, then for all $X \subset A$ such that $a \in X$, $b \notin N_G(X)$, we have $|N_G(X)| > |X|$. In particular, since $H' \supseteq G$, then for our X we have $a \in X$ and $b \notin N_G(X)$ and thus $|N_{H'}(X)| \geq |N_G(X)| > |X|$, in contradiction to the definition of X. □

Non-Totally Ordered Graphs Have a Zero Coefficient

Recall that the only monomials which may appear in $BPM_n^\star$ are those corresponding to graphs $G \in HVC_n$. Combining this characterization with those obtained using the Eulerian lattice of matching-covered graphs, we get:

Lemma 6.2.9. *Let $n > 1$ and let $G \in HVC_n$. If G is **not totally ordered**, then G has a surplus edge.*

Proof. Let $A = \{a_1, \ldots, a_n\}$, $B = \{b_1, \ldots, b_n\}$ be two sets, and let $G = (A \cup B, E) \in HVC_n$, such that G is not totally ordered. Thus, there exist two vertices $a_i, a_j \in A$ such that:

$$N(a_i) \nsubseteq N(a_j) \quad \wedge \quad N(a_j) \nsubseteq N(a_i) \quad \wedge \quad |N(a_i)| \geq |N(a_j)|$$

We will show that $\forall b_k \in (N(a_j) \setminus N(a_i)) : (a_i, b_k)$ is a surplus edge for G. Let $b_k \in N(a_j) \setminus N(a_i)$, $b_m \in N(a_i) \setminus N(a_j)$. Since $G \in HVC_n$, every edge of G is covered by some graph $K \in HV_n$, and in particular so are (a_i, b_m), (a_j, b_k). Thus, there exist $X_{i,m}, X_{j,k} \subseteq A$, $Y_{i,m}, Y_{j,k} \subseteq B$ such that:

$$|X_{i,m}| + |Y_{i,m}| = n+1, \qquad |X_{j,k}| + |Y_{j,k}| = n+1$$
$$(X_{i,m} \times Y_{i,m}) \subseteq E(G) \qquad (X_{j,k} \times Y_{j,k}) \subseteq E(G)$$

and furthermore, $a_i \in X_{i,m}$, $b_m \in Y_{i,m}$, $a_j \in X_{j,k}$ and $b_k \in Y_{j,k}$. Assume towards a contradiction that (a_i, b_k) is not a surplus edge for G. Then, there exists $X \subset A$ such that $a_i \in X$, $b_k \notin N(X)$ and $|N(X)| \leq |X|$.

Since $a_i \in X$ then $N(X) \supseteq N(a_i)$ and in particular $|N(X)| \geq |N(a_i)|$. Furthermore, observe that $X \cap X_{j,k} = \varnothing$, since otherwise $b_k \in N(X)$, in contradiction to the definition of X. Thus, we have $n - |X_{j,k}| \geq |X|$. Moreover, recall that by the definition of a_i and a_j, we have $|N(a_i)| \geq |N(a_j)|$. Since the edge (a_j, b_k) is covered by $K_{X_{j,k}, Y_{j,k}}$, then $N(a_j) \supseteq Y_{j,k}$. Lastly, by the definition of X, $|N(X)| \leq |X|$. Putting all the above inequalities together, we have:

$$n - |X_{j,k}| \geq |X| \geq |N(X)| \geq |N(a_i)| \geq |N(a_j)| \geq |Y_{j,k}|$$

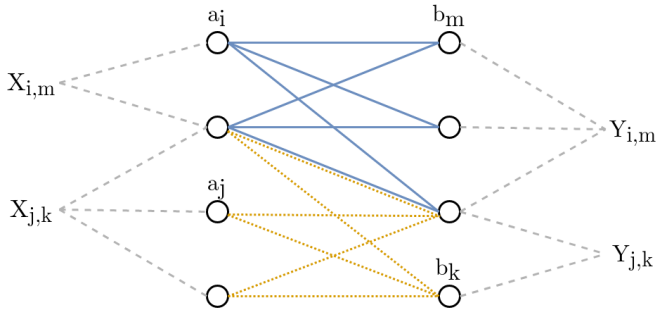

Figure 6.1: A graph $G \in HVC_4$, which is not totally ordered.
The edge (a_i, b_m) is covered by $K_{X_{i,m},Y_{i,m}}$, and (a_j, b_k) is covered by $K_{X_{j,k},Y_{j,k}}$.

Therefore $|X_{j,k}| + |Y_{j,k}| \leq n$, in contradiction to the fact that $|X_{j,k}| + |Y_{j,k}| = n + 1$. □

Corollary 6.2.9.1. *Let $n > 1$ and let $G \subseteq K_{n,n}$. If G is **not totally ordered**, then $a^\star_G = 0$.*

Strictly Totally Ordered Graphs Have a Non-Zero Coefficient

Lemma 6.2.10. *Let $n > 1$ and let $G \subseteq K_{n,n}$ be **strictly totally ordered**. Then:*

$$a^\star_G = (-1)^{n+1}$$

Proof. Let $G = (A \cup B, E) \subseteq K_{n,n}$ be a graph, where $A = \{a_1, \ldots, a_n\}$ and $B = \{b_1, \ldots, b_n\}$. The edges of G are given by: $\forall i \in [n] : N_G(a_i) = \{b_1, \ldots, b_i\}$. Observe that G is **strictly totally ordered**, since $N(a_n) \supsetneq N(a_{n-1}) \supsetneq \ldots N(a_1) \supsetneq \varnothing$.

For every $k \in [n]$, denote $A_k = \{a_k, \ldots, a_n\} \subseteq A$, $B_k = \{b_1, \ldots, b_k\} \subseteq B$. By the definition of G, $\forall k \in [n]$, $K_{A_k,B_k} \subseteq G$. We now show that for any $K_{X,Y} \in HV_n$ such that $K_{X,Y} \subseteq G$ and $|Y| = k$, we necessarily have $X = A_k$ and $Y = B_k$.

Assume towards a contradiction this is not the case. Let $K_{X,Y} \in HV_n$ such that $K_{X,Y} \subseteq G$, $|Y| = k$ and $Y \neq B_k$ or $X \neq A_k$. If $Y = B_k$, then for any $a_i \in X$ where $a_i \notin A_k$ (i.e., $i < k$), the edge $(a_i, b_1) \notin E(G)$ – a contradiction. Otherwise, let $Y \neq B_k$ and let $j > k$ the maximal index such that $y_j \in Y$. By the definition of G, $\bigcap_{b \in Y} N_G(b) = N_G(y_j) = A_j$. Since $K_{X,Y} \subseteq G$, then in particular $X \subseteq A_j$, and therefore $|Y| + |X| \leq k + |A_j| = n$, in contradiction to the fact that $K_{X,Y} \in HV_n$.

Thus, the only Hall violators appearing in G are the set:

$$\mathcal{H} := \{H \in HV_n : H \subseteq G\} = \{K_{A_k,B_k} : k \in [n]\}$$

The union of all graphs in $\mathcal{H}$ is exactly G, therefore $G \in HVC_n$. However, recall that $BPM_n^\star$ is a graph cover function for the set HV_n, and thus by arithmetizing the formula representing the function (recall Proposition 5.1.2), we get that:

$$BPM_n^\star(x_{1,1},\ldots,x_{n,n}) = \sum_{G\in HVC_n}\left(\sum_{\substack{\emptyset\neq S\subseteq HVC_n\\ \bar{E}(S)=E(G)}}(-1)^{|S|+1}\right)\prod_{(i,j)\in E(G)}x_{i,j}$$

Observe that the *only* set of graphs $S \subseteq HVC_n$ whose union is equal to G is the set $\mathcal{H}$ itself, since by omitting any K_{A_k,B_k} we will fail to cover the edge $(a_k, b_k) \in E(G)$. Thus $a_G^\star = (-1)^{|\mathcal{H}|+1} = (-1)^{n+1}$, as required.

Lastly, we observe that **any** strictly totally ordered graph $G' \subseteq K_{n,n}$ is equivalent, up to permutations over each bipartition, to G (and therefore has the same coefficient). This equivalence can be achieved by sorting the vertices of each bipartition by the cardinality of their neighbour sets, where the left vertices are sorted in ascending order, and the right vertices in descending order. □

This concludes the proof of Theorem 2.

6.3 Counting the Monomials of $BPM_n^\star$

Using Theorem 2, we now deduce the following asymptotically tight bound on the number of monomials appearing in $BPM^\star$.

Corollary 6.3.0.1. *Let $n > 1$. The number of monomials in $BPM_n^\star$ satisfies:*

$$(n!)^2 \leq |mon(BPM_n^\star)| < (n+2)^{2n+2}$$

And in particular:

$$\log_2\left(|mon(BPM_n^\star)|\right) = 2n\log_2 n \pm \mathcal{O}(n) = \Theta(n\log n)$$

Proof. Let $n > 1$. For the lower bound, let G be a strictly totally ordered graph. By Lemma 6.2.10, *all* strictly totally ordered graphs, and in particlar G, have $a_G^\star = (-1)^{n+1}$. However, since no two right or left vertices of G have the same set of neighbours, any pair of permutations

over the left and right bipartitions yields a new strictly totally ordered graph $\tilde{G} \cong G$, thus completing the lower bound.

For the upper bound, let $U = \{u_1, \ldots, u_{n+1}\}$, $V = \{v_1, \ldots, v_{n+1}\}$ be two sets. Denote by C_n the set of all graphs $G \subseteq K_{n,n}$ that are *totally ordered*. We begin by showing that:

$$|C_n| = \sum_{k=1}^{n+1} \left((k-1)! \cdot \left\{ {n+1 \atop k} \right\} \right)^2$$

Where the notation $\left\{ {n \atop k} \right\}$ refers to the Stirling number of the second kind. To prove the equality, let us explicitly construct the set C_n as follows; for every $1 \le k \le n+1$, let:

$$U = U_1 \uplus U_2 \cdots \uplus U_k \qquad V = V_1 \uplus V_2 \cdots \uplus V_k$$

be partitions of U, V, respectively, into k non-empty subsets, where without loss of generality $u_{n+1} \in U_k$ and $v_{n+1} \in V_k$. Then, for every $\pi, \tau \in S_{k-1}$, consider the graph $G \in C_n$, whose edges are given by:

$$\forall i \in [k-1] : \ \forall u \in U_{\pi(i)} : \ N_G(u) = V_{\tau(1)} \uplus \cdots \uplus V_{\tau(i)}$$

Recall that the number of partitions of n elements into k non-empty subsets is given by $\left\{ {n \atop k} \right\}$, the Stirling number of the second kind. Thus by the above construction, the cardinality of the set C_n satisfies:

$$|C_n| = \sum_{k=1}^{n+1} \underbrace{((k-1)!)^2}_{\text{Choosing } \pi,\tau} \cdot \underbrace{\left\{ {n+1 \atop k} \right\}^2}_{\text{Partitioning } U,V}$$

Therefore:

$$\begin{aligned} |mon(BPM_n^\star)| \le |C_n| &= \sum_{k=1}^{n+1} \left((k-1)! \cdot \left\{ {n+1 \atop k} \right\} \right)^2 \\ &\le \left(\sum_{k=1}^{n+1} k! \cdot \left\{ {n+1 \atop k} \right\} \right)^2 = (F_{n+1})^2 \end{aligned}$$

Where F_n denotes the n'th Fubini number. We now use the upper bound [Mez19]: $\forall n \ge 1 : F_n < (n+1)^n$, thereby concluding the proof. □

Is the Totally Ordered Condition Necessary?

Since $BPM_n^\star$ is a graph cover function for HV_n, the only monomials which *may* appear in $BPM_n^\star$ are those corresponding to graphs $G \in HVC_n$ – i.e., graphs covered by Hall violators. Clearly the number of Hall violators is $\Omega(2^{2n})$, however, one might wonder about a corresponding *upper bound* for the number of graphs *covered* by Hall violators. In particular, could we

perhaps have derived as strong an asymptotic bound as the one yielded by the *totally ordered* condition (Definition 6.1.4), by simply bounding the size of the set HVC_n? The following proposition shows that this is not the case, namely, there are (asymptotically) many more graphs which are *covered* by Hall violators:

Proposition 6.3.1. *Let $n > 1$. Then:*

$$\log_2(|HVC_n|) \geq \left\lfloor \frac{n}{2} \right\rfloor \left(\left\lceil \frac{n}{2} \right\rceil + 1 \right) \geq \frac{n^2}{4} - 1$$

Proof. Let $n > 1$ and without loss of generality assume that $n = 2k$ where $k \in \mathbb{N}^+$. Let A,B be two sets such that $|A| = |B| = n$. The lower bound follows by constructing a graph $G = (A \cup B, E_G) \in HVC_n$ where $|E(G)| = n^2 - n/2(n/2+1)$, such that $\{H \supseteq G\} \subseteq HVC_n$. First, partition each bipartition A,B into two sets, as follows:

$$A = (X \cup Y): \quad X = \{a_1, \ldots, a_k\} \quad , \quad Y = \{a_{k+1}, \ldots, a_{2k}\}$$
$$B = (U \cup V): \quad U = \{b_1, \ldots, b_{k-1}\} \, , \quad V = \{b_k, \ldots, b_{2k}\}$$

The edges of G are formed by connecting all edges between X and B, and all edges between Y and U, thus: $E(G) = (X \times B) \cup (Y \times U)$. Observe that $G \in HVC_n$, since it can be covered by taking k copies of $K_{x,B}$, one for each $x \in X$, and taking another $k-1$ copies of $K_{A,u}$, one for each $u \in U$.

Any missing edge $(y, v) \notin E(G)$ (where $y \in Y$ and $v \in V$) can be covered by $K_{X \cup \{y\}, U \cup \{v\}}$ (the complete bipartite graph connecting $X \cup \{y\}$ and $U \cup \{v\}$). Observe that $K_{X \cup \{y\}, U \cup \{v\}} \in HV_n$, since $|X \cup U \cup \{y, v\}| = n + 1$. Thus $\{H \supseteq G\} \subseteq HVC_n$, as required. □

6.4 Corollaries of Theorem 2

Communication Matrix Rank

Consider the following communication problem. Given an input graph $G \subseteq K_{n,n}$, its edges are distributed between two parties, Alice and Bob, according to some arbitrary fixed partition. The parties' task is to devise a communication protocol to determine whether G contains a bipartite perfect matching. This communication problem is known as the *"2-Party Communication Problem of Bipartite Perfect Matching"*.

Clearly, the rank of the associated communication matrix for this problem is at least exponential in n (e.g., using a fooling set argument). Interestingly, the compact representation of $BPM^{\star}_n$ given by Theorem 2 allows us to deduce that the rank of the communication matrix is, in fact, *at most* exponential in $n \log n$.

Corollary 6.4.0.1. *Let M be the communication matrix for the 2-party communication problem of bipartite perfect matching. The rank of M over the Reals is bounded by:*

$$rank_{\mathbb{R}}(M) \leq (n+2)^{2n+2} = 2^{\mathcal{O}(n \log n)}$$

Proof. Let M be the aforementioned communication matrix, and let $\bar{M} = J - M$, where J is the all-ones matrix. The polynomial $BPM_n^\star$ induces an (at most) $|mon(BPM_n^\star)|$-rank decomposition of $\bar{M}$, since each monomial is a rank-1 matrix (see, e.g., [NW95]). However, $rank(M) \leq rank(\bar{M}) + 1$, and by Corollary 6.3.0.1, $|mon(BPM_n^\star)| < (n+2)^{2n+2}$, thus concluding the proof. □

Lower Bound for OR Decision Trees

Much in the same way that the multilinear polynomial representing BPM_n allowed us to derive query complexity lower bounds for AND decision trees, the multilinear polynomial representing $BPM_n^\star$ can be used to obtain similar lower bounds against OR decision trees. The proof is very similar to that of Lemma 5.4.1, but differs in several key steps, thus we provide it for completeness.

Lemma 6.4.1. *Let $f : \{0,1\}^n \to \{0,1\}$ be a boolean function. Then:*

$$D^{OR}(f) \geq \log_3(|mon(f^\star)|)$$

Where $f^\star$ is the dual function of f.

Proof. Let T be an OR-decision tree computing f and denote $d = depth(T)$. Let $\mathcal{P}$ be the set of all root to 0-leaf paths in T. For any $P \in \mathcal{P}$, the indicator function for the path is given by the following multilinear polynomial:

$$\mathbb{1}_P(x_1, \ldots, x_n) = \left(\prod_{\neg OR(S) \in P} \left(1 - \sum_{\varnothing \neq S \subseteq [n]} (-1)^{|S|+1} \prod_{i \in S} x_i \right) \right) \left(\prod_{OR(S) \in P} \left(\sum_{\varnothing \neq S \subseteq [n]} (-1)^{|S|+1} \prod_{i \in S} x_i \right) \right)$$

By the definition of the dual function, we can construct $f^\star$ by summing the indicators for all paths $P \in \mathcal{P}$, where the inputs to each indicator are the negated input bits:

$$\begin{aligned} f^\star(x_1, \ldots, x_n) &= \sum_{P \in \mathcal{P}} \mathbb{1}_P(1 - x_1, \ldots, 1 - x_n) \\ &= \sum_{P \in \mathcal{P}} \left(\prod_{OR(S) \in P} \left(1 - \prod_{i \in S} x_i \right) \right) \left(\prod_{\neg OR(S) \in P} \left(\prod_{i \in S} x_i \right) \right) \end{aligned}$$

Therefore, each path P making k right turns contributes at most 2^k monomials to $f^\star$. In a binary tree of depth d there are at most $\binom{d}{k}$ paths making exactly k right turns (i.e., by selecting the position in the path at which the right turns are made). Thus, we have:

$$|mon(f^\star)| \leq \sum_{k=0}^{d} \binom{d}{k} 2^k = 3^d$$ □

Applying the aforementioned lemma to BPM_n, we obtain:

Corollary 6.4.1.1. *The depth of any OR decision tree computing BPM_n is at least:*

$$D^{OR}(BPM_n) \geq 2\log_3(n!)$$

6.5 Additional Coefficients of the Dual Polynomial

Theorem 2 offers a characterization of $BPM_n^\star$ in terms of *totally ordered* and *strictly totally ordered* graphs. The theorem states that only *totally ordered* graphs may exhibit non-zero coefficients, and that all *strictly totally ordered* indeed have non-zero coefficients. For graphs that are totally ordered but not strictly so, the situation is more complex [1]. The following proposition shows that for any $n > 2$, there exist graphs which are totally ordered but not strictly so, whose dual coefficient is 1, 0 and even $(n-2)^2$.

Proposition 6.5.1. *Let $n > 2$. There exist graphs $G \subseteq K_{n,n}$ which are totally ordered but not strictly so, such that:*

$$a.\ a_G^\star = 0, \qquad b.\ a_G^\star = 1, \qquad c.\ a_G^\star = (n-2)^2$$

Proof. Let $n > 2$ and let $A = \{a_1, \ldots, a_n\}$, $B = \{b_1, \ldots, b_n\}$ be two sets. Denote:

$$A_{n-1} = \{a_1, \ldots, a_{n-1}\}, \quad B_{n-1} = \{b_1, \ldots, b_{n-1}\}$$

For the case $a_G^\star = 0$, consider any totally ordered graph such that $G \in (MC_n \setminus \{K_{n,n}\})$. For example, let $G = (A \cup B, E)$ such that $\forall i \in [n-1]:\ N_G(a_i) = B$ and $N_G(a_n) = \{b_1, b_2\}$. G is totally ordered, since $N_G(a_1) \supseteq N_G(a_2) \supseteq \cdots \supseteq N_G(a_n)$. However, we also have that $G \in MC_n$, therefore by Lemma 6.2.2, $a_G^\star = 0$.

[1] An additional analysis exclusively for graphs containing a perfect matching can be found in Appendix A

For the case $a^\star_G = 1$, consider any graph $G \in HV_n$. Observe that G is both totally ordered and $G \in HVC_n$. Furthermore G contains a *single* Hall violator graph (itself), and is therefore a minterm of $BPM^\star_n$, and so $a^\star_G = 1$.

Lastly, for the case $a^\star_G = (n-2)^2$, consider the graph $G = K_{A_{n-1},B_{n-1}}$. Using Theorem 5.2.3, the set of matching-covered graphs containing G, which we will denote by $\mathcal{H}$, can be partitioned into three sets $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2 \cup \mathcal{H}_3$, as follows:

$$\mathcal{H}_1 = \{G \cup \{(a_n, b_n)\}\}$$
$$\mathcal{H}_2 = \{G \cup \{(a_n, b_n)\} \cup (U \times \{b_n\}) \cup (\{a_n\} \times V) : \varnothing \neq U \subseteq A_{n-1},\ \varnothing \neq V \subseteq B_{n-1}\}$$
$$\mathcal{H}_3 = \{G \cup (U \times \{b_n\}) \cup (\{a_n\} \times V) : U \subseteq A_{n-1},\ V \subseteq B_{n-1},\ |U| \geq 2,\ |V| \geq 2\}$$

By Lemma 6.2.1, the dual coefficient of G is given by:

$$a^\star_G = (-1)^{|E(G)|+1} \sum_{\substack{H \supseteq G \\ H \in MC_n}} (-1)^{\chi(H)} = -\sum_{H \in \mathcal{H}} (-1)^{|E(H) \setminus E(G)| + |C(H)|}$$

For the single graph $H \in \mathcal{H}_1$, $|E(H) \setminus E(G)| = 1$, $|C(H)| = 2$, thus contributing 1 to the sum. For each $H \in \mathcal{H}_2$, $|C(H)| = 1$, thus $\mathcal{H}_2$'s contribution to the sum is:

$$\sum_{i=1}^{n-1} \sum_{j=1}^{n-1} \binom{n-1}{i} \binom{n-1}{j} (-1)^{i+j+1} = -1$$

Lastly, for each $H \in \mathcal{H}_3$, $|C(H)| = 1$. Thus $\mathcal{H}_3$'s contribution to the sum is:

$$\sum_{i=2}^{n-1} \sum_{j=2}^{n-1} \binom{n-1}{i} \binom{n-1}{j} (-1)^{i+j} = (n-2)^2$$

Summing up all the contributions, we get $a^\star_G = (n-2)^2$, thus concluding the proof.

□

Bibliography

[BS94] Louis J Billera and Aravamuthan Sarangarajan.
The combinatorics of permutation polytopes.
In *Formal power series and algebraic combinatorics*, volume 24, pages 1–23, 1994.

[DNO19] Shahar Dobzinski, Noam Nisan, and Sigal Oren.
Economic efficiency requires interaction.
Games and Economic Behavior, 118:589–608, 2019.

[ER64] Paul Erdős and Alfred Rényi.
On random matrices.
Magyar Tud. Akad. Mat. Kutató Int. Közl, 8(455-461):1964, 1964.

[Fri85] Alan M Frieze.
Limit distribution for the existence of hamiltonian cycles in random bipartite graphs.
European Journal of Combinatorics, 6(4):327–334, 1985.

[Grü13] Branko Grünbaum.
Convex polytopes, volume 221.
Springer Science & Business Media, 2013.

[Het64] Gábor Hetyei.
Rectangular configurations which can be covered by 2×1 rectangles.
Pécsi Tan. Foisk. Közl, 8:351–367, 1964.

[HK73] John E Hopcroft and Richard M Karp.
An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs.
SIAM Journal on computing, 2(4):225–231, 1973.

[KSS84] Jeff Kahn, Michael Saks, and Dean Sturtevant.
A topological approach to evasiveness.
Combinatorica, 4(4):297–306, 1984.

[Mez19] Istvan Mezo.
Combinatorics and number theory of counting sequences.
Chapman and Hall/CRC, 2019.

[ML19] Sagnik Mukhopadhyay and Bruno Loff.
Lifting theorems for equality.
In *STACS 2019*, 2019.

[Nis19] Noam Nisan.
The demand query model for bipartite matching.
arXiv preprint arXiv:1906.04213, 2019.

[NSS08] Ilan Newman, Michael Saks, and Mario Szegedy.
Decision trees - unpublished manuscript.
2008.

[NW95] Noam Nisan and Avi Wigderson.
On rank vs. communication complexity.
Combinatorica, 15(4):557–565, 1995.

[O'D14] Ryan O'Donnell.
Analysis of boolean functions.
Cambridge University Press, 2014.

[PL86] M.D. Plummer and L. Lovász.
Matching Theory.
North-Holland Mathematics Studies. Elsevier Science, 1986.

[RV75] Ronald L Rivest and Jean Vuillemin.
A generalization and proof of the Aanderaa-Rosenberg conjecture.
In *Proceedings of the seventh annual ACM symposium on Theory of computing*, pages 6–11. ACM, 1975.

[Sta11] Richard P. Stanley.
Enumerative Combinatorics: Volume 1.
Cambridge University Press, New York, NY, USA, 2nd edition, 2011.

[Yao88] Andrew Chi-Chih Yao.
Monotone bipartite graph properties are evasive.
SIAM Journal on Computing, 17(3):517–520, 1988.

APPENDIX A

Graphs with a Perfect Matching

In this section, we restrict our attention to graphs *containing a perfect matching*, which appear in the dual polynomial $BPM_n^\star$. By Theorem 2, the only graphs appearing in the dual polynomial are those which are "totally ordered". However, by nature of having a perfect matching, a more precise characterization of their structure can be obtained.

Given a graph G with a perfect matching, we consider the graph G', formed by the union of *all* perfect matchings of G. In this section, we show that if the monomial corresponding to G appears $BPM_n^\star$, then the following conditions (and perhaps others) must hold. First, all the connected components of G' must be complete bipartite graphs. Furthermore, for any edge in G connecting two such components, all the edges between the components' corresponding bipartitions must appear.

Lemma A.0.1. *Let $n > 1$ and let $G \subseteq K_{n,n}$, where $G \notin MC_n$ and $PM(G) \neq \varnothing$. Denote by G' the union of all the perfect matchings of G. If G' has a connected component which is not a complete bipartite graph, then $a_G^\star = 0$.*

Proof. Let $G = (A \cup B, E) \notin MC_n$, where $PM(G) \neq \varnothing$ and denote by G' the union of all perfect matchings of G. Let C be a connected component of G' which is not a complete bipartite graph. Let $(a,b) \in (A \cap V(C)) \times (B \cap V(C))$ be an edge such that $(a,b) \notin E(C)$. We will show that (a,b) is a wildcard edge for G. Therefore, let $H \in MC_n$ be a graph such that $H \supseteq G \cup \{(a,b)\}$), and denote by $\tilde{H}$ the connected component of H containing C.

Observe that $\tilde{H} - V(C)$ contains a perfect matching (in particular, any of the perfect matchings induced by the components C_i of G' which are contained in $\tilde{H}$). Thus, $\tilde{H}$ has a bipartite ear decomposition of the form: $\tilde{H} = C + P_1 + \cdots + P_q$, where there exists a path $P_i = (a,b)$ (since the vertices a, b were present in C). Therefore $\tilde{H} \setminus \{(a,b)\}$ also has a bipartite ear decomposition: $\tilde{H} \setminus \{(a,b)\} = C + P_1 + \cdots + P_{i-1} + P_{i+1} + \cdots + P_q$, and by Theorem 5.2.5, $\tilde{H} \setminus \{(a,b)\}$ is elementary. Thus $(H \setminus \{(a,b)\}) \in MC_n$ and the proof follows by Lemma 6.2.6. □

Lemma A.0.2. *Let $n > 1$ and let $G = (A \cup B, E) \subseteq K_{n,n}$. Denote G' the union of G's perfect matchings. If all the following conditions hold:*

1. *$G \in HVC_n$ and $PM(G) \neq \varnothing$.*

2. *All the connected components of G' are complete bipartite graphs.*

3. *There exist $C_1 = (A_1 \cup B_1, E_1)$, $C_2 = (A_2 \cup B_2, E_2)$, where $C_1, C_2 \in C(G')$, such that:*

$$\varnothing \subsetneq ((A_1 \times B_2) \cap E(G)) \subsetneq (A_1 \times B_2)$$

Then $a^\star_G = 0$.

Proof. Let G be a graph satisfying the above conditions, and let G', C_1, C_2 be the graphs described above. Denote $C(G') = \{C_1, \ldots, C_t\}$, where $\forall i \in [t] : \ C_i = (A_i \cup B_i, E_i)$. Hereafter, we use the notation $C_i \rightsquigarrow C_j$ to denote an edge $(u, v) \in (A_i \times B_j)$.

First, since $G \in HVC_n$ and G has a perfect matching, then G is connected. Let $(a, b), (u, v) \in (A_1 \times B_2)$ be two edges, such that $(a, b) \notin E(G)$ and $(u, v) \in E(G)$. We will show that (a, b) is a wildcard edge of G. Let $H \in MC_n$ be a graph such that $H \supseteq G \cup \{(a, b)\}$). We will show that $H' = H \setminus \{(a, b)\}$ is elementary, thus by Lemma 6.2.6, $a^\star_G = 0$. Let $(x, y) \in E(H')$. To show that H' is elementary, by Theorem 5.2.3 it is sufficient to exhibit a perfect matching of H' containing (x, y).

Clearly, if $\exists i \in [t] : (x, y) \in E(C_i)$ then since C_i is elementary, $C_i - x - y$ has a perfect matching, which can be extended to a perfect matching of H' by adding a single perfect matching for each $C_j \in (C(G') \setminus C_i)$.

Otherwise, denote by C_i, C_j the components for which $x \in C_i$, $y \in C_j$. We begin by showing that H has a directed cycle $\bar{C} = C_i \rightsquigarrow C_j \rightsquigarrow \cdots \rightsquigarrow C_i$ containing (x, y). Since $H \in MC_n$, every edge of H participates in a perfect matching, and in particular so does (x, y). Let M be a perfect matching of H involving (x, y). Since $C_i - x$ is unbalanced, there must be some edge $C_k \rightsquigarrow C_i$ in M. Iteratively applying the same argument to C_k and then to the component connected to it, we eventually gather a directed cycle $\bar{C} \in E(H)$ composed of edges of M, where $(x, y) \in \bar{C}$.

Lastly, we use $\bar{C}$ to construct a perfect matching of H' containing (x, y). First, if $(a, b) \in \bar{C}$, then replace (a, b) with (u, v). Now, construct a perfect matching $\bar{M}$ as follows:

1. For each $C_k \notin \bar{C}$, take a single perfect matching over C_k.

2. For each edge $(a_k, b_m) \in \bar{C}$, match a_k and b_m.

3. For each $C_k \in \bar{C}$, denote $a_k \in A_k$, $b_k \in B_k$ the vertices of C_k appearing in $\bar{C}$. By Theorem 5.2.3, $C_k - a_k - b_k$ has at least one perfect matching (or is empty if $C_k = K_2$), which we add to $\bar{M}$.

□

APPENDIX B

The Polynomial of $BPM_3^\star$

$$
\begin{aligned}
BPM_3^\star(x) = {} & x_{1,1}x_{1,2}x_{1,3} + x_{1,1}x_{2,1}x_{3,1} + x_{2,1}x_{2,2}x_{2,3} + x_{1,2}x_{2,2}x_{3,2} + x_{1,3}x_{2,3}x_{3,3} + x_{3,1}x_{3,2}x_{3,3} \\
& + x_{1,1}x_{1,2}x_{2,1}x_{2,2} + x_{1,1}x_{1,3}x_{2,1}x_{2,3} + x_{1,2}x_{1,3}x_{2,2}x_{2,3} + x_{1,1}x_{1,2}x_{3,1}x_{3,2} + x_{1,1}x_{1,3}x_{3,1}x_{3,3} \\
& + x_{1,2}x_{1,3}x_{3,2}x_{3,3} + x_{2,1}x_{2,2}x_{3,1}x_{3,2} + x_{2,1}x_{2,3}x_{3,1}x_{3,3} + x_{2,2}x_{2,3}x_{3,2}x_{3,3} - x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2} \\
& - x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,3} - x_{1,1}x_{1,2}x_{1,3}x_{2,2}x_{2,3} - x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{3,1} - x_{1,1}x_{1,2}x_{2,1}x_{2,2}x_{2,3} \\
& - x_{1,1}x_{1,2}x_{1,3}x_{2,2}x_{3,2} - x_{1,1}x_{1,3}x_{2,1}x_{2,2}x_{2,3} - x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3} - x_{1,1}x_{1,2}x_{1,3}x_{2,3}x_{3,3} \\
& - x_{1,1}x_{1,2}x_{2,1}x_{2,2}x_{3,1} - x_{1,1}x_{1,2}x_{2,1}x_{2,2}x_{3,2} - x_{1,1}x_{1,3}x_{2,1}x_{2,3}x_{3,1} - x_{1,1}x_{1,2}x_{1,3}x_{3,1}x_{3,2} \\
& - x_{1,1}x_{1,2}x_{1,3}x_{3,1}x_{3,3} - x_{1,1}x_{1,2}x_{1,3}x_{3,2}x_{3,3} - x_{1,1}x_{1,3}x_{2,1}x_{2,3}x_{3,3} - x_{1,2}x_{1,3}x_{2,2}x_{2,3}x_{3,2} \\
& - x_{1,2}x_{1,3}x_{2,2}x_{2,3}x_{3,3} - x_{1,1}x_{1,2}x_{2,1}x_{3,1}x_{3,2} - x_{1,1}x_{1,2}x_{2,2}x_{3,1}x_{3,2} - x_{1,1}x_{2,1}x_{2,2}x_{2,3}x_{3,1} \\
& - x_{1,1}x_{1,3}x_{2,1}x_{3,1}x_{3,3} - x_{1,2}x_{2,1}x_{2,2}x_{2,3}x_{3,2} - x_{1,1}x_{1,3}x_{2,3}x_{3,1}x_{3,3} - x_{1,2}x_{1,3}x_{2,2}x_{3,2}x_{3,3} \\
& - x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,3} - x_{1,2}x_{1,3}x_{2,3}x_{3,2}x_{3,3} - x_{1,1}x_{2,1}x_{2,2}x_{3,1}x_{3,2} - x_{1,2}x_{2,1}x_{2,2}x_{3,1}x_{3,2} \\
& - x_{1,1}x_{1,2}x_{3,1}x_{3,2}x_{3,3} - x_{1,1}x_{2,1}x_{2,3}x_{3,1}x_{3,3} - x_{1,1}x_{1,3}x_{3,1}x_{3,2}x_{3,3} - x_{1,3}x_{2,1}x_{2,3}x_{3,1}x_{3,3} \\
& - x_{1,2}x_{1,3}x_{3,1}x_{3,2}x_{3,3} - x_{1,2}x_{2,2}x_{2,3}x_{3,2}x_{3,3} - x_{1,3}x_{2,2}x_{2,3}x_{3,2}x_{3,3} - x_{1,1}x_{2,1}x_{3,1}x_{3,2}x_{3,3} \\
& - x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2} - x_{1,2}x_{2,2}x_{3,1}x_{3,2}x_{3,3} - x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,3} - x_{2,1}x_{2,2}x_{2,3}x_{3,2}x_{3,3} \\
& - x_{1,3}x_{2,3}x_{3,1}x_{3,2}x_{3,3} - x_{2,1}x_{2,2}x_{3,1}x_{3,2}x_{3,3} - x_{2,1}x_{2,3}x_{3,1}x_{3,2}x_{3,3} - x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3} \\
& + 2x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3} + x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{3,1} + x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,3}x_{3,1} \\
& + x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{3,2} + x_{1,1}x_{1,2}x_{1,3}x_{2,2}x_{2,3}x_{3,2} + x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,3}x_{3,3} \\
& + x_{1,1}x_{1,2}x_{1,3}x_{2,2}x_{2,3}x_{3,3} + x_{1,1}x_{1,2}x_{2,1}x_{2,2}x_{2,3}x_{3,1} + x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{3,1}x_{3,2} \\
& + x_{1,1}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,1} + x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{3,1}x_{3,3} + x_{1,1}x_{1,2}x_{2,1}x_{2,2}x_{2,3}x_{3,2} \\
& + x_{1,1}x_{1,2}x_{1,3}x_{2,2}x_{3,1}x_{3,2} + x_{1,1}x_{1,2}x_{1,3}x_{2,3}x_{3,1}x_{3,3} + x_{1,1}x_{1,2}x_{1,3}x_{2,2}x_{3,2}x_{3,3} \\
& + x_{1,1}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,3} + x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,2} + x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,3} \\
& + x_{1,1}x_{1,2}x_{1,3}x_{2,3}x_{3,2}x_{3,3} + 2x_{1,1}x_{1,2}x_{2,1}x_{2,2}x_{3,1}x_{3,2} + 2x_{1,1}x_{1,3}x_{2,1}x_{2,3}x_{3,1}x_{3,3} \\
& + 2x_{1,1}x_{1,2}x_{1,3}x_{3,1}x_{3,2}x_{3,3} + 2x_{1,2}x_{1,3}x_{2,2}x_{2,3}x_{3,2}x_{3,3} + x_{1,1}x_{1,2}x_{2,1}x_{3,1}x_{3,2}x_{3,3} \\
& + x_{1,1}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2} + x_{1,1}x_{1,3}x_{2,1}x_{3,1}x_{3,2}x_{3,3} + x_{1,1}x_{1,2}x_{2,2}x_{3,1}x_{3,2}x_{3,3} \\
& + x_{1,2}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2} + x_{1,1}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,3} + x_{1,2}x_{2,1}x_{2,2}x_{2,3}x_{3,2}x_{3,3} \\
& + x_{1,1}x_{1,3}x_{2,3}x_{3,1}x_{3,2}x_{3,3} + x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,3} + x_{1,2}x_{1,3}x_{2,2}x_{3,1}x_{3,2}x_{3,3} \\
& + x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,2}x_{3,3} + x_{1,2}x_{1,3}x_{2,3}x_{3,1}x_{3,2}x_{3,3} + x_{1,1}x_{2,1}x_{2,2}x_{3,1}x_{3,2}x_{3,3} \\
& + x_{1,2}x_{2,1}x_{2,2}x_{3,1}x_{3,2}x_{3,3} + x_{1,1}x_{2,1}x_{2,3}x_{3,1}x_{3,2}x_{3,3} + x_{1,2}x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3} \\
& + x_{1,3}x_{2,1}x_{2,3}x_{3,1}x_{3,2}x_{3,3} + x_{1,3}x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3} + 2x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3} \\
& - x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,1} - x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,2} - x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,3} \\
& - x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{3,1}x_{3,2} - x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,3}x_{3,1}x_{3,3} - x_{1,1}x_{1,2}x_{1,3}x_{2,2}x_{2,3}x_{3,2}x_{3,3} \\
& - x_{1,1}x_{1,2}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2} - x_{1,1}x_{1,2}x_{1,3}x_{3,1}x_{3,2}x_{3,3} - x_{1,1}x_{1,2}x_{1,3}x_{2,2}x_{3,1}x_{3,2}x_{3,3} \\
& - x_{1,1}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,3} - x_{1,1}x_{1,2}x_{1,3}x_{2,3}x_{3,1}x_{3,2}x_{3,3} - x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,2}x_{3,3} \\
& - x_{1,1}x_{1,2}x_{2,1}x_{2,2}x_{3,1}x_{3,2}x_{3,3} - x_{1,1}x_{1,3}x_{2,1}x_{2,3}x_{3,1}x_{3,2}x_{3,3} - x_{1,2}x_{1,3}x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3} \\
& - x_{1,1}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3} - x_{1,2}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3} - x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3} \\
& + x_{1,1}x_{1,2}x_{1,3}x_{2,1}x_{2,2}x_{2,3}x_{3,1}x_{3,2}x_{3,3}
\end{aligned}
$$

The Monomials of $BPM_4^\star$

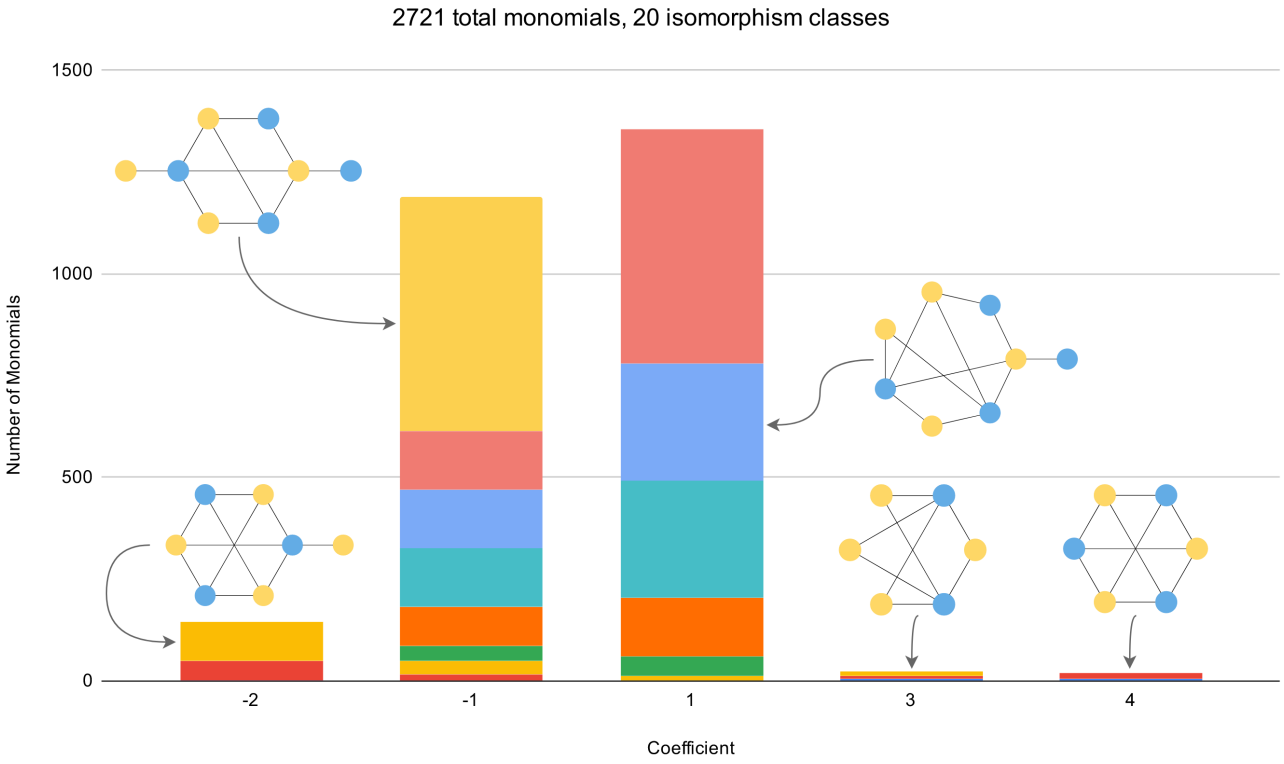

Figure C.1: The monomials of $BPM_4^\star$, grouped by their coefficient.
For each coefficient, different colours indicate isomorphism classes.

The Lattice of Matching-Covered Graphs, for $n = 3$

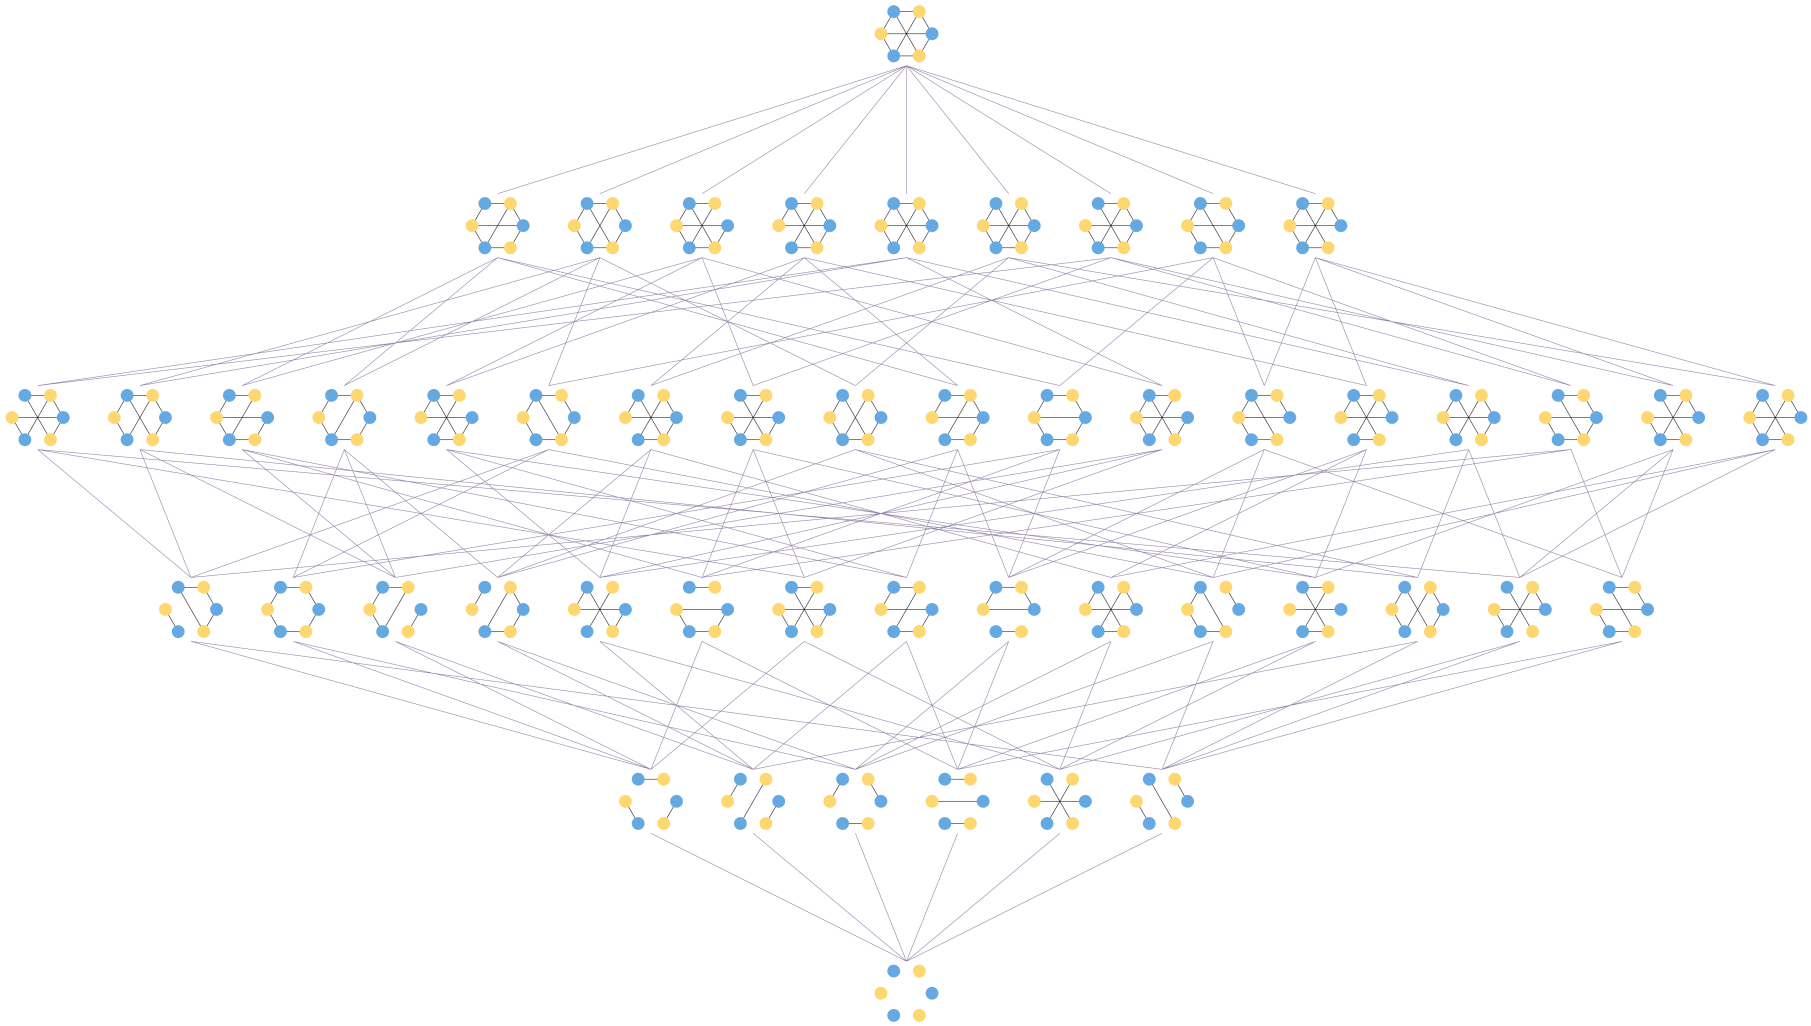

Figure D.1: The Lattice $\mathscr{P} = (MC_3 \cup \{\hat{0}\}, \subseteq)$, which is isomorphic to the face lattice of the Birkhoff Polytope B_3