Algebraic Representations of Unique Bipartite Perfect Matching

Gal Beniamini ✉
The Hebrew University of Jerusalem, Israel

Abstract

We obtain complete characterizations of the Unique Bipartite Perfect Matching function, and of its Boolean dual, using multilinear polynomials over the reals. Building on previous results [2, 3], we show that, surprisingly, the dual description is *sparse* and has *low* ℓ_1*-norm* – only exponential in $\Theta(n \log n)$, and this result extends even to other families of matching-related functions. Our approach relies on the Möbius numbers in the matching-covered lattice, and a key ingredient in our proof is Möbius' inversion formula.

These polynomial representations yield complexity-theoretic results. For instance, we show that unique bipartite matching is *evasive* for classical decision trees, and *nearly evasive* even for generalized query models. We also obtain a tight $\Theta(n \log n)$ bound on the log-rank of the associated two-party communication task.

2012 ACM Subject Classification Mathematics of computing → Matchings and factors; Theory of computation → Communication complexity; Theory of computation → Oracles and decision trees

Keywords and phrases Bipartite Perfect Matching, Boolean Functions, Partially Ordered Sets

Digital Object Identifier 10.4230/LIPIcs.MFCS.2022.16

1 Introduction

A perfect matching in a graph is a subset of edges spanning the graph, no two of which are incident to the same vertex. In this paper we consider the *decision problem* of unique bipartite matching: the input is a balanced bipartite graph over $2n$ vertices, and the goal is to determine whether the graph contains a *unique* perfect matching. This problem can be naturally cast as a Boolean function.

▶ **Definition.** *The unique bipartite perfect matching function* $\mathrm{UBPM_n} : \{0,1\}^{n^2} \to \{0,1\}$ *is*

$$\mathrm{UBPM_n}(x_{1,1}, \ldots, x_{n,n}) = \begin{cases} 1 & \{(i,j) : x_{i,j} = 1\} \text{ has a unique perfect matching} \\ 0 & \text{otherwise.} \end{cases}$$

The complexity of $\mathrm{UBPM_n}$ is closely related to that of $\mathrm{BPM_n}$ – the problem in which we drop the uniqueness condition and simply ask whether a bipartite graph *contains* a perfect matching. Both $\mathrm{BPM_n}$ and $\mathrm{UBPM_n}$ are known to lie in **P**, due to a classical result by Edmonds [9]. However, despite their close connection, not all known algorithmic results extend from one problem to another. For instance, $\mathrm{UBPM_n}$ was shown by Kozen, Vazirani and Vazirani to be in **NC** [19] (see also [14]), and no such result is known for $\mathrm{BPM_n}$. Lovász showed that $\mathrm{BPM_n}$ is in **RNC** [21], and the current best-known *deterministic* parallel algorithm is due to Fenner, Gurjar and Thierauf [10], placing the problem in **Quasi-NC**. Determining the membership of bipartite perfect matching in **NC** remains one of the main open problems in parallelizability.

Our main results in this paper are the complete characterizations of both $\mathrm{UBPM_n}$ and its dual function, by means of polynomials. These characterizations leverage a deep connection to the polynomial representations of $\mathrm{BPM_n}$, obtained in [3, 2], and it is our hope that they can be used to further our understanding of the connection between the two. To present our results we require some notation. We say that a bipartite graph is *matching-covered* if

every edge of the graph participates in some perfect matching. For a graph G we denote its *cyclomatic number*, a topological quantity, by $\chi(G) = e(G) - v(G) + c(G)$. The set of all perfect matchings of G is denoted $\mathrm{PM}(G)$, and the cardinality of this set is denoted $\mathrm{per}(G)$ (the permanent of G). Under these notations, our first theorem is the following closed-form description of the *unique* real multilinear polynomial representing $\mathrm{UBPM_n}$.

▶ **Theorem 1** (The Unique Bipartite Perfect Matching Polynomial)**.**

$$\mathrm{UBPM_n}(x_{1,1},\ldots,x_{n,n}) = \sum_{G\subseteq K_{n,n}} c_G \prod_{(i,j)\in E(G)} x_{i,j}$$

where

$$c_G = \begin{cases} (-1)^{\chi(G)}\,\mathrm{per}(G) & G \text{ is matching-covered} \\ 0 & \text{otherwise.} \end{cases}$$

The polynomial appearing in Theorem 1 bears a striking resemblance to the representation of $\mathrm{BPM_n}$, the only difference being the multiplicative $\mathrm{per}(G)$ appearing in each term of $\mathrm{UBPM_n}$. This is a direct result of the connection between the two functions and the *matching-covered lattice*, hereafter $\mathcal{L}_n$, which is formed by all matching-covered graphs of order $2n$, ordered with respect to the subgraph relation. Billera and Sarangarajan [5] proved that $\mathcal{L}_n$ is isomorphic to the face lattice of the Birkhoff Polytope $\mathbf{B}_n$. Consequently, this combinatorial lattice is Eulerian, and its Möbius function is particularly well-behaved – a fact which we rely on indirectly throughout this paper. In [3], it was shown that $\mathrm{BPM_n}$ is intimately related to the matching-covered lattice: every such graph corresponds to a monomial, and their coefficients are given by Möbius numbers. Our proof of Theorem 1 extends this connection by leveraging Möbius Inversion Formula, and in fact allows us to derive the polynomial representation for *any* indicator function over $\mathcal{L}_n$ (including, for instance, $\mathrm{BPM_n}$), while also simplifying somewhat parts of the original proof.

Theorem 1 yields information-theoretic lower bounds. For example, $\mathrm{UBPM_n}$ has full total degree and is thus *evasive*, i.e., any decision tree computing it must have full depth, n^2. Unlike its analogue $\mathrm{BPM_n}$, which is a *monotone* bipartite graph property and thus known to be evasive [17], the *unique* perfect matching function is *not monotone*, and for such functions evasiveness is not guaranteed (see e.g. [23]). We also obtain lower bounds against generalized families of decision trees, whose internal nodes are labeled by arbitrary parity functions (XOR-DT), or conjunctions (AND-DT), over subsets of the inputs bits.

▶ **Corollary.** *For classical, parity, and conjunction trees, the following lower bounds hold:*

$$\mathrm{D}(\mathrm{UBPM_n}) = n^2, \quad \mathrm{D}^{\mathrm{XOR}}(\mathrm{UBPM_n}) \geq \left(\tfrac{1}{2} - o(1)\right) n^2 \quad \textit{and} \quad \mathrm{D}^{\mathrm{AND}}(\mathrm{UBPM_n}) \geq (\log_3 2)n^2 - o(1).$$

In the second part of this paper we consider the Boolean dual function $\mathrm{UBPM_n^\star}$, which is obtained by flipping all the input and output bits (or formally, $\mathrm{UBPM_n^\star}(x_{1,1},\ldots,x_{n,n}) = 1 - \mathrm{UBPM_n}(1-x_{1,1},\ldots,1-x_{n,n})$). By construction, this is the indicator over all bipartite graphs whose *complement* does *not* contain a unique perfect matching. Our second result is a complete characterization of $\mathrm{UBPM_n^\star}$ as a real multilinear polynomial. This description relies *heavily* on the that of $\mathrm{BPM_n^\star}$ – which is the dual of the bipartite perfect matching function $\mathrm{BPM_n}$. The polynomial representation of the latter dual was obtained in a series of papers [3, 2], and is omitted here for brevity.

▶ **Theorem 2** (The Dual Polynomial of Unique Bipartite Perfect Matching)**.**

$$\mathrm{UBPM}^{\star}_{\mathrm{n}}(x_{1,1},\ldots,x_{n,n}) = \sum_{G\subseteq K_{n,n}} c^{\star}_{G} \prod_{(i,j)\in E(G)} x_{i,j}$$

where

$$c^{\star}_{G} = \mathrm{per}(G)\cdot a^{\star}_{G} + \sum_{M\notin \mathrm{PM}(G)} (-1)^{|E(M)\setminus E(G)|}\cdot a^{\star}_{G\cup M}$$

and $a^{\star}_{G}$ *denotes the coefficient of* G *in* $\mathrm{BPM}^{\star}_{\mathrm{n}}$*.*

Theorem 2 expresses the coefficient of every graph G as an alternating sum over coefficients of $\mathrm{BPM}^{\star}_{\mathrm{n}}$, corresponding exactly to those graphs formed by adjoining a single perfect matching to G. This suffices in order to *inherit* the main structural result of [2] regarding $\mathrm{BPM}^{\star}_{\mathrm{n}}$: the ℓ_1-norm of $\mathrm{UBPM}^{\star}_{\mathrm{n}}$, i.e., the norm of the coefficient vector of the representing polynomial, is *very small* – only exponential in $\Theta(n\log n)$, and this is tight.

▶ **Corollary.** *The dual polynomial is sparse and its coefficients are small. Explicitly,*

$$\log \left\|\mathrm{UBPM}^{\star}_{\mathrm{n}}\right\|_1 = \Theta(n\log n).$$

The low norm of the dual yields algorithmic results for the unique-bipartite-matching problem, and for related matching problems. For instance, through the approximation scheme of [2, 29], it allows one to obtain a low-degree polynomial *approximation* of the unique bipartite matching function over the hypercube (i.e., "approximate degree"), which holds even for *exponentially small* error. The same ℓ_1-norm bound also directly extends to the spectral norm of $\mathrm{UBPM}_{\mathrm{n}}$,[1] which is a well-studied quantity in analysis of Boolean functions.

Finally, we consider the two-party deterministic communication complexity of unique bipartite matching. The input is a graph $G\subseteq K_{n,n}$, whose edges are distributed among two parties according to *any arbitrary* and *fixed* partition. The sparse polynomial representation of $\mathrm{UBPM}^{\star}_{\mathrm{n}}$ allows us to deduce that the log-rank of the communication matrix, for *any* of the above communication tasks, is bounded by only $\mathcal{O}(n\log n)$, and we prove that this is tight.[2] We remark that, while we show that unique matching has low log-rank, not much is known regarding its *deterministic communication complexity*. For the monotone variant $\mathrm{BPM}_{\mathrm{n}}$, known algorithms (e.g. [15]) can be translated into protocols using only $\widetilde{\mathcal{O}}(n^{3/2})$ bits [25]. However, it is currently not known how to convert algorithms for $\mathrm{UBPM}_{\mathrm{n}}$ (such as [11, 12]), into protocols using even $\mathcal{O}(n^{2-\varepsilon})$ bits, for any $\varepsilon>0$. Determining the deterministic communication complexity of $\mathrm{UBPM}_{\mathrm{n}}$ is thus left as an open problem.

2 Preliminaries and Notation

2.1 Boolean Functions and Polynomials

Every Boolean function $f:\{0,1\}^n\to\{0,1\}$ can be *uniquely* represented by a multilinear polynomial $p\in\mathbb{R}[x_1,\ldots,x_n]$ (see e.g. [27]), where f and p agree on all Boolean inputs $\{0,1\}^n$. The family of subsets corresponding to monomials in this polynomial representation

[1] The spectra of any function and its dual are identical up to sign, and the $\{0,1\}$-polynomial ℓ_1-norm is always trivially at least as large as the $\{\pm1\}$-representation ("Fourier") ℓ_1-norm.

[2] In fact, our results hold even for a certain $\wedge$-*lifted* and dualised version of this problem.

(i.e., whose coefficient does not vanish) is denoted by $\mathrm{mon}(f)$. The cardinality of $\mathrm{mon}(f)$ is known as the *sparsity* of f, and the maximal cardinality of any $S \in \mathrm{mon}(f)$ is known as the *total degree* of f, hereafter $\deg(f)$. The ℓ_1-norm of f is the norm of its representing polynomial's coefficient vector, namely:

$$\|f\|_1 \stackrel{\text{def}}{=} \left\|(a_S)_{S\subseteq[n]}\right\|_1, \text{ where } f \text{ is } \{0,1\}\text{-represented by } p(x_1,\ldots,x_n) = \sum_{S\subseteq[n]} a_S \prod_{i\in S} x_i.$$

Given a Boolean function $f : \{0,1\}^n \to \{0,1\}$, it is often useful to consider the transformation in which we *invert* all the input and output bits. This process produces a new Boolean function $f^\star$, known as the Boolean dual.

▶ **Definition 3.** *Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function. The **Boolean dual** of f is the function $f^\star : \{0,1\}^n \to \{0,1\}$ where the symbols 0 are 1 are interchanged. Formally,*

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

The polynomial representations of a Boolean function f and its dual $f^\star$ can differ substantially (for example $\mathrm{AND_n}^\star = \mathrm{OR_n}$, and while the former is represented by a single monomial, the latter consists of $2^n - 1$ monomials). However, since f and $f^\star$ are obtained by affine transformations of one another, they share many properties. For example, their Fourier spectra are identical [27], up to sign. Moreover, they have the same approximate degree [2] for any error ε, and the ranks of their associated communication matrices (see proceeding subsections) are identical up to an additive constant (of 1).

2.2 Graphs

We use the standard notation for quantities relating to graphs. In particular, the sets of vertices, edges and connected components of a graph are denoted by $V(G)$, $E(G)$ and $C(G)$, and their cardinalities are denoted $v(G)$, $e(G)$ and $c(G)$, respectively. A less common measure appearing in this paper is the cyclomatic number $\chi(G)$, a topological quantity.

▶ **Definition 4.** *Let G be a graph. The **cyclomatic number** of G is defined by:*

$$\chi(G) = e(G) - v(G) + c(G).$$

A *matching* in a graph $G \subseteq K_{n,n}$ is a collection of edges sharing no vertices, and said matchings are called *perfect* if they contain exactly n edges (i.e., every vertex in the graph is incident to precisely one edge in the matching). The set of *all* perfect matchings denoted by $\mathrm{PM}(G)$. For any graph $G \subseteq K_{n,n}$, we define the *permanent* $\mathrm{per}(G)$ and the *determinant* $\det(G)$ as the application of these two functions to the biadjacency matrix of G, noting that $\mathrm{per}(G)$ counts the number of perfect matchings in G.

Perfect matchings and the graphs formed by unions thereof play a central role in this paper. A graph $G \subseteq K_{n,n}$ is called **matching-covered** if and only if every edge of G participates in some perfect matching. Matching-covered graphs have interesting combinatorial properties. For example, this is precisely the family of all graphs admitting a bipartite ear decomposition (similar to the ear decomposition of 2-edge-connected graphs). This family had previously appeared extensively in the literature, and in particular had been studied at length by Lovász and Plummer [28], and by Hetyei [13]. Hereafter, we denote the set of all such graphs by

$$\mathrm{MC_n} = \Big\{G \subseteq K_{n,n} : G \text{ is matching-covered}\Big\}.$$

All graphs in this paper are *balanced bipartite graphs*, over the fixed vertex set of the complete bipartite graph $K_{n,n}$. Consequently, we use the notation $G \subseteq H$ to indicate inclusion over the edges, and similarly $G \cup H$ is the graph whose edges are $E(G) \cup E(H)$. Lastly, many of the Boolean functions appearing in this paper are defined over subgraphs of $K_{n,n}$, where every input bit is associated with a single edge. For such functions, the notation $f(G)$, where $G \subseteq K_{n,n}$, corresponds to this mapping.

2.3 Communication Complexity

In this paper we consider the *two-party deterministic communication model*. For a comprehensive textbook on the topic, we refer the reader to [20]. The **deterministic communication complexity** of f, hereafter $\mathrm{D}^{\mathrm{CC}}(f)$, is the *least number of bits communicated by a protocol computing f, on the worst-case input.* Any (*unpartitioned*) Boolean function $f : \{0,1\}^n \to \{0,1\}$ naturally gives rise to a *family* of associated two-party communication tasks: one corresponding to each possible partition of the input bits between the two parties. The **deterministic communication complexity of a Boolean function** is then defined as follows.

▶ **Definition 5.** *The deterministic communication complexity of $f : \{0,1\}^n \to \{0,1\}$ is defined*

$$\mathrm{D}^{\mathrm{CC}}(f) \stackrel{def}{=} \max_{S \sqcup \bar{S} = [n]} \mathrm{D}^{\mathrm{CC}}\left(f_S(x,y)\right)$$

where $\mathrm{D}^{\mathrm{CC}}\left(f_S(x,y)\right)$ is the deterministic communication complexity of the two-party Boolean function $f_S(x,y) : \{0,1\}^{|S|} \times \{0,1\}^{|\bar{S}|} \to \{0,1\}$, representing f under the partition $S \sqcup \bar{S}$.

For any two-party Boolean function, let us also define the following two useful objects.

▶ **Definition 6.** *Let $f : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}$. The **communication matrix** of f is*

$$M_f \in \mathbb{R}^{\{0,1\}^m \times \{0,1\}^n}, \text{ where } \forall (x,y) \in \{0,1\}^m \times \{0,1\}^n : \ M_f(x,y) = f(x,y).$$

▶ **Definition 7.** *Let $f : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}$ be a two-party function. We say that $S \subseteq \{0,1\}^m \times \{0,1\}^n$ is a **fooling set** for f if and only if:*

$$S \subseteq f^{-1}(1), \text{ and } \forall (x_1,y_1) \neq (x_2,y_2) \in S : \ \{(x_1,y_2),(x_2,y_1)\} \cap f^{-1}(0) \neq \emptyset.$$

The log of the rank of M_f *over the reals* (sometimes referred to as the "log-rank of f") is intimately related to the communication complexity of f. A classical theorem due to Mehlhorn and Schmidt [24] states that $\mathrm{D}^{\mathrm{CC}}(f) \geq \log_2 \operatorname{rank} M_f$, and these two quantities are famously conjectured to be polynomially related [22]. As for the fooling set, it is well known that $\mathrm{D}^{\mathrm{CC}}(f) \geq \log_2 \mathrm{fs}(f)$ for any two-party function f [20], where $\mathrm{fs}(f)$ is the maximum size of a fooling set. This bound was extended by Dietzfelbinger, Hromkovič and Schnitger [8], who showed that in fact $\log \mathrm{fs}(f) \leq 2 \log \operatorname{rank} f + 2$.

2.4 Posets, Lattices and Möbius Functions

Partially ordered sets (hereafter, **posets**) are defined by a tuple $\mathcal{P} = (P, \leq)$, where P is the element set, and $\leq$ is the order relation (which is reflexive, antisymmetric and transitive). For any two elements $x, y \in P$, the notation $[x,y] \stackrel{\text{def}}{=} \{z \in P : x \leq z \leq y\}$ denotes the *interval* from x to y. A *combinatorial lattice* is a poset satisfying two additional conditions: every two

elements have a least upper bound (a "join"), and a greatest lower bound (a "meet"). The *face lattice of a polytope* is a combinatorial lattice whose elements correspond to the faces of a polytope, ordered by the subset relation. Such a lattice is *bounded* – it has a unique bottom element (the empty face $\hat{0}$), and a unique top element (the polytope itself), and it is also *graded*, meaning that the length of all maximal chains between any two elements x, y are identical (in other words, the elements can be *ranked*).

Partially ordered sets come equipped with an important function known as the **Möbius function**. The Möbius function of a poset is the inverse, with respect to convolution, of its zeta function $\zeta(x, y) = \mathbb{1}\{x < y\}$. For information on incidence algebra and the Möbius function, we refer the reader to [30].

▶ **Definition 8** (Möbius Function for Posets). *Let $\mathcal{P} = (P, \leq)$ be a finite poset. The Möbius function $\mu_{\mathcal{P}} : P \times P \to \mathbb{R}$ of $\mathcal{P}$ is defined*

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

The Möbius Inversion Formula allows one to relate two functions defined on a poset $\mathcal{P}$, where one function is a downwards closed sum of another, by means of the Möbius function. This can be seen as a generalization of its number-theoretic analogue (as indeed the Möbius function of number theory arises in this manner from the *divisibility poset*).

▶ **Theorem 9** (Möbius Inversion Formula, see [30]). *Let $\mathcal{P} = (P, \leq)$ be a finite poset and let $f, h : P \to \mathbb{F}$ be two functions, where $\mathbb{F}$ is a field. Then:*

$$\forall x \in P: \ h(x) = \sum_{y \leq x} f(y) \quad \Longleftrightarrow \quad \forall x \in P: \ f(x) = \sum_{y \leq x} h(y)\mu_{\mathcal{P}}(y, x).$$

3 The Unique Perfect Matching Polynomial

Our main object of study is the unique bipartite matching function.

▶ **Definition 10.** *The Unique Bipartite Perfect Matching function is defined*

$$\mathrm{UBPM_n}(x_{1,1}, \ldots, x_{n,n}) = \begin{cases} 1 & \{(i,j) : x_{i,j} = 1\} \subseteq K_{n,n} \text{ has a unique P.M.} \\ 0 & \text{otherwise.} \end{cases}$$

The unique multilinear representation of $\mathrm{UBPM_n}$ is characterized in the following Theorem.

▶ **Theorem 1.** *The unique polynomial $\mathrm{UBPM_n} : \{0,1\}^{n^2} \to \{0,1\}$ is given by*

$$\mathrm{UBPM_n}\left(x_{1,1}, \ldots, x_{n,n}\right) = \sum_{G \in \mathrm{MC_n}} (-1)^{\chi(G)} \ \mathrm{per}(G) \prod_{(i,j) \in E(G)} x_{i,j}.$$

Proof. The proof is centered around the combinatorial *lattice of matching-covered graphs*,

$$\mathcal{L}_n = \left(\mathrm{MC_n} \cup \{\hat{0}\}, \subseteq\right), \text{ where } \hat{0} \text{ is the graph with } 2n \text{ isolated vertices}$$

where the order relation for this lattice is *containment over the edge set*, i.e., $G \supseteq H \iff E(G) \supseteq E(H)$. Let us consider the following two functions $f : \mathcal{L}_n \to \{0,1\}$ and $h : \mathcal{L}_n \to \mathbb{Z}$ on the lattice, which are the restrictions of $\mathrm{UBPM_n}$ and of the Permanent function, respectively:

$$\forall G \in (\mathrm{MC_n} \cup \{\hat{0}\}): \ f(G) = \mathrm{UBPM_n}(G) = \begin{cases} 1 & G \in PM(K_{n,n}) \\ 0 & \text{otherwise} \end{cases}$$

$$h(G) = \mathrm{per}(G) = \#\text{Perfect Matchings in } G.$$

These two functions are intimately related. Indeed, for any element G of the lattice, one can compute $h(G)$ by taking the sum $f(H)$ over all H in the downwards closed interval $[\hat{0}, G]$. Therefore, by an application of Möbius' Inversion Formula (Theorem 9) to the matching-covered lattice, we obtain:

$$\forall G \in \mathcal{L}_n : \; h(G) = \sum_{G \supseteq H \in \mathcal{L}_n} f(H) \iff \forall G \in \mathcal{L}_n : \; f(G) = \sum_{G \supseteq H \in \mathcal{L}_n} \mu(H, G) h(H)$$

where $\mu : \mathcal{L}_n \to \mathbb{Z}$ is the Möbius function of the lattice $\mathcal{L}_n$. A well known result due to Billera and Sarangarajan [5] states that $\mathcal{L}_n$ is isomorphic to the *face lattice* of the Birkhoff Polytope B_n, which is the convex hull of all $n \times n$ permutation matrices. Consequently, $\mathcal{L}_n$ is an Eulerian lattice – and its Möbius function μ is can be directly computed (see e.g. [30]), as follows:

$$\forall G, H \in \mathcal{L}_n, \; H \subseteq G : \quad \mu(H, G) = (-1)^{\mathrm{rank}(G) - \mathrm{rank}(H)}$$

where $\mathrm{rank}(x)$ denotes the maximal length of a chain from $\hat{0}$ to x (equivalently, $\mathrm{rank}(x) = dim(f_x) + 1$, where f_x is the face of B_n corresponding to the lattice element x). In [3] it was shown that the rank of every graph G in the matching-covered lattice is exactly $\chi(G) + 1$, where $\chi(G) = e(G) - v(G) + c(G)$ is the cyclomatic number, a topological quantity. Recalling our prior application of Möbius inversion, we obtain the following set of identities (note that the bottom element can be omitted, as $\mathrm{per}(\hat{0})$ is zero):

$$\forall G \in \mathcal{L}_n : \; (-1)^{\chi(G)} \sum_{G \supseteq H \in \mathrm{MC_n}} (-1)^{\chi(H)} \mathrm{per}(H) = \begin{cases} 1 & G \in \mathrm{PM}(K_{n,n}) \\ 0 & \text{otherwise.} \end{cases}$$

To conclude the proof, let us consider the following real multilinear polynomial, wherein we assign weight $(-1)^{\chi G} \mathrm{per}(G)$ to every matching-covered graph:

$$p(x_{1,1}, \ldots, x_{n,n}) = \sum_{G \in \mathrm{MC_n}} (-1)^{\chi(G)} \mathrm{per}(G) \prod_{(i,j) \in E(G)} x_{i,j}.$$

Let $G \subseteq K_{n,n}$ and observe that, by construction:

$$p(G) = \sum_{H \in \mathrm{MC_n}} (-1)^{\chi(H)} \mathrm{per}(H) \cdot \mathbb{1}\{E(H) \subseteq E(G)\} = \sum_{G \supseteq H \in \mathrm{MC_n}} (-1)^{\chi(H)} \mathrm{per}(H).$$

It remains to show that p "agrees" with $\mathrm{UBPM_n}$ on all inputs. It is not hard to see that it suffices to show this claim only for matching-covered graphs, since given any $G \subseteq K_{n,n}$ which is *not* matching-covered, one may consider the graph G' formed by the union of all perfect matching in G (in other words, the maximal matching-covered graph contained in G). By construction, we have $p(G') = p(G)$, and by definition, $\mathrm{UBPM_n}(G) = \mathrm{UBPM_n}(G')$ – thus, hereafter we consider only inputs $G \in \mathrm{MC_n}$. First, let us check the two trivial cases; the empty graph, and a single matching:

$$p(\hat{0}) = 0, \text{ and } p(M) = (-1)^{\chi(M)} = (-1)^{n-2n+n} = 1 \quad \forall M \in PM(K_{n,n}).$$

Finally, for any matching-covered graph G containing *more than a single matching*, i.e. $G \in \mathrm{MC_n}$ such that $G \notin PM(K_{n,n})$, it holds that:

$$p(G) = \sum_{G \supseteq H \in \mathrm{MC_n}} (-1)^{\chi(H)} \mathrm{per}(H) = 0$$

where the last equality follows from the identities obtained through Möbius' Inversion Formula. Thus, $p(x_{1,1}, \ldots, x_{n,n})$ agrees with $\mathrm{UBPM_n}$ everywhere, and is its *unique* representation. ◀

3.1 Indicators on the Matching-Covered Lattice

We remark that the proof of Theorem 1 readily extends, through the same analysis using Möbius inversion, to any arbitrary *indicator function* over the matching-covered lattice. For any set $S \subseteq \mathrm{MC_n}$, let $I_S : \{0,1\}^{n^2} \to \{0,1\}$ be the Boolean function

$$\forall G \subseteq K_{n,n} : \; I_S(G) = \mathbb{1}\left\{ H \in S \text{ where } H = \bigcup_{M \in \mathrm{PM}(G)} M \right\}.$$

Then, the multilinear polynomial representing I_S is given by

$$I_S(x_{1,1}, \ldots, x_{n,n}) = \sum_{G \in \mathrm{MC_n}} \left((-1)^{\chi(G)} \sum_{H \in [\hat{0}, G] \cap S} (-1)^{\chi(H)+1} \right) \prod_{(i,j) \in E(G)} x_{i,j}.$$

3.2 Evasiveness and Generalized Decision Trees

The characterization of $\mathrm{UBPM_n}$ as a multilinear polynomial can be used to derive several complexity-theoretic corollaries. Firstly, this polynomial has *full total degree over* $\mathbb{R}$ and thus (see e.g. [6]):

▶ **Corollary 11.** $\mathrm{UBPM_n}$ *is evasive, i.e., any decision computing it has full depth,* n^2.

Let us remark that, contrary to its counterpart $\mathrm{BPM_n}$ which is a *monotone* bipartite graph property and thus known to be evasive [31], the *unique* matching function is *not monotone* and for such functions evasiveness is not guaranteed (see [23] for one such example). Theorem 1 can be also used to derive strong bounds (near evasiveness) versus larger classes of decision trees, for example trees whose internal nodes are labeled by arbitrary conjunctions of the input bits (hereafter AND-DT), and by arbitrary parity functions (XOR-DT). It is known [3] that the depth of any AND-DT computing a Boolean function f is at least $\log_3 |\operatorname{mon}(f)|$. Applying this to $\mathrm{UBPM_n}$ and recalling that asymptotically almost all balanced bipartite graphs are matching-covered ([3]), we have:

▶ **Corollary 12.** *Any* AND-DT *computing* $\mathrm{UBPM_n}$ *has depth at least* $(\log_3 2) \cdot n^2 - o_n(1)$.

As for parity decision trees, it is well known that the depth of any such tree is bounded by the total degree of its unique representing polynomial, over $\mathbb{F}_2$ (see [27, 3]). Noting that $\mathrm{per}(G) \equiv \det(G) \pmod 2$, we may write the $\mathbb{F}_2$-polynomial representation of $\mathrm{UBPM_n}$ as follows

$$\mathrm{UBPM_n}\,(x_{1,1}, \ldots, x_{n,n}) = \sum_{\substack{G \in \mathrm{MC_n} \\ \det(G) \equiv 1 \pmod 2}} \prod_{(i,j) \in E(G)} x_{i,j}.$$

Clearly this polynomial does not have full degree for any $n > 1$, as $\mathrm{per}(K_{n,n})$ is $n! \equiv 0 \pmod 2$[3]. Nevertheless, we claim that its $\mathbb{F}_2$-degree is at most a constant factor away from full. Observe that its monomials constitute precisely of all graphs that are both matching-covered, and whose biadjacency matrices are invertible over $\mathbb{F}_2$, i.e., are elements of the group $\mathrm{GL}_n(\mathbb{F}_2)$. However, asymptotically almost all graphs are matching-covered, and by a standard counting argument, the order of $\mathrm{GL}_n(\mathbb{F}_2)$ satisfies

[3] It is well known ([27]) that for any function $f : \{0,1\}^n \to \{0,1\}$, $\deg_2(f) = n \iff |f^{-1}(1)| \equiv 1 \pmod 2$. Therefore we obtain that the number of graphs $G \subseteq K_{n,n}$ containing a *unique perfect matching* is even, for any $n > 1$.

$$\Pr_{A\sim M_n(\mathbb{F}_2)}[A \in \mathrm{GL}_n(\mathbb{F}_2)] = \left(\tfrac{1}{2};\tfrac{1}{2}\right)_\infty \pm o_n(1)$$

where $\left(\frac{1}{2};\frac{1}{2}\right)_\infty \approx 0.28878$ is a Pochhammer symbol. Thus by a standard Chernoff argument, there exists a matching-covered graph with odd determinant and at least $\frac{1}{2}n^2 - o_n(1)$ edges.

▶ **Corollary 13.** $\mathrm{D}^{\mathrm{XOR}}(\mathrm{UBPM_n}) \geq \deg_2(\mathrm{UBPM_n}) \geq \left(\frac{1}{2} - o_n(1)\right) n^2$

4 The Dual Polynomial

In this section we consider the Boolean dual function (Definition 3) of $\mathrm{UBPM_n}$.

▶ **Definition 14.** *The function* $\mathrm{UBPM_n^\star} : \{0,1\}^{n^2} \to \{0,1\}$ *is defined*

$$\mathrm{UBPM_n^\star}(x_{1,1},\ldots,x_{n,n}) = \begin{cases} 1 & \{(i,j) : x_{i,j} = 0\} \subseteq K_{n,n} \text{ does } \underline{\text{not}} \text{ have a unique P.M.} \\ 0 & \text{otherwise.} \end{cases}$$

In what follows, we provide a full characterization of polynomial representing $\mathrm{UBPM_n^\star}$. This description relies heavily on the that of another dual function – $\mathrm{BPM_n^\star}$ – which is the dual of the bipartite perfect matching function $\mathrm{BPM_n}$ (which is defined identically to $\mathrm{UBPM_n}$, but without the *uniqueness* condition). The polynomial representation of $\mathrm{BPM_n^\star}$ was obtained in a series of papers [3, 2]. Its monomials correspond to a family of graphs called "*totally ordered graphs*", and their coefficients are can be computed through a normal-form block decomposition of the aforementioned graphs. The full details are presented in [2], and are omitted here for brevity. In what follows, it suffices for us to denote

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

Under this notation, our characterization of $\mathrm{UBPM_n^\star}$ is the following.

▶ **Theorem 2.** *The unique polynomial representation of* $\mathrm{UBPM_n^\star} : \{0,1\}^{n^2} \to \{0,1\}$ *is*

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

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

$$c_G^\star = \mathrm{per}(G)\cdot a_G^\star + \sum_{M\notin \mathrm{PM}(G)} (-1)^{|E(M)\setminus E(G)|}\cdot a_{G\cup M}^\star.$$

Proof. The polynomial representing $\mathrm{UBPM_n^\star}$ can be expressed using $\mathrm{UBPM_n}$, via duality:

$$\mathrm{UBPM_n^\star}(x_{1,1},\ldots,x_{n,n}) = 1 - \mathrm{UBPM_n}(1-x_{1,1},\ldots,1-x_{n,n}).$$

Substituting the characterization of Theorem 1 and expanding, we deduce that the coefficient of every graph $G \subseteq K_{n,n}$ in $\mathrm{UBPM_n}$ is:

$$c_G^\star = (-1)^{e(G)+1} \sum_{G\subseteq H\in \mathrm{MC_n}} (-1)^{\chi(H)}\,\mathrm{per}(H).$$

Writing $\mathrm{per}(H) = \sum_{M\in\mathrm{PM}(K_{n,n})} \mathbb{1}\{M\subseteq H\}$ and exchanging order of summation,

$$c_G^\star = (-1)^{e(G)+1} \sum_{M\in\mathrm{PM}(K_{n,n})} \sum_{G\subseteq H\in \mathrm{MC_n}} (-1)^{\chi(H)} \mathbb{1}\{M\subseteq H\}.$$

There are two possible cases in the above summation over all perfect matchings; either the matching is present in G, or it is not. Clearly every matching-covered graph containing G also contains any matching of G, so in the former case we get a contribution of $(-1)^{e(G)+1}\operatorname{per}(G)\cdot\sum_{G\subseteq H\in\mathrm{MC_n}}(-1)^{\chi(H)}$. As for the latter case, observe that for every $M\notin\mathrm{PM}(G)$, the set of matching-covered graphs containing G and M is exactly all matching-covered graphs containing $G\cup M$. Finally, we recall [3] that the coefficient of any graph $G\subseteq K_{n,n}$ in $\mathrm{BPM}^\star_\mathrm{n}$ is given by:

$$a^\star_G=(-1)^{e(G)+1}\sum_{G\subseteq H\in\mathrm{MC_n}}(-1)^{\chi(H)}.$$

Putting the two together and simplifying, we obtain:

$$c^\star_G=\operatorname{per}(G)\cdot a^\star_G+\sum_{M\notin\mathrm{PM}(G)}(-1)^{|E(M)\setminus E(G)|}\cdot a^\star_{G\cup M}.$$ ◀

4.1 Corollary: The ℓ_1-norm of $\mathrm{UBPM}^\star_\mathrm{n}$

One immediately corollary of Theorem 2 is the following fact: the multilinear polynomial representing $\mathrm{UBPM}^\star_\mathrm{n}$ has *very low ℓ_1-norm* – i.e., it has few monomials, and the coefficient of every such monomial is not too large. A similar bound had previous been attained for $\mathrm{BPM}^\star_\mathrm{n}$ in [2], which we heavily rely on for our proof.

▶ **Corollary 15.** *The ℓ_1-norm of $\mathrm{UBPM}^\star_\mathrm{n}$ is bounded only by $\log_2\|\mathrm{UBPM}^\star_\mathrm{n}\|_1=\Theta(n\log n)$.*

Proof. For the upper bound, we rely heavily on Theorem 2 and on the ℓ_1-norm of $\mathrm{BPM}^\star_\mathrm{n}$ obtained in [2]. In the latter, it was shown that every coefficient in $\mathrm{BPM}^\star_\mathrm{n}$ has magnitude at most 2^{2n}, and thus using the characterization of Theorem 2, the coefficient of any graph G satisfies

$$\log_2|c^\star_G|\le\log_2\left(\operatorname{per}(G)\cdot 2^{2n}+(n!-\operatorname{per}(G))\cdot 2^{2n}\right)\le n\log_2 n+n\log_2(4/e)+\Theta(\log n).$$

It remains to bound the *sparsity* of $\mathrm{UBPM}^\star_\mathrm{n}$. To this end, consider the graphs whose coefficients do not vanish in $\mathrm{BPM}^\star_\mathrm{n}$, and let us take a "ball" around every such graph $G\in\operatorname{mon}(\mathrm{BPM}^\star_\mathrm{n})$, as follows:

$$B(G)=\left\{H\subseteq K_{n,n}:\exists M\in\mathrm{PM}(K_{n,n})\text{ such that }E(H)\cup E(M)=E(G)\right\}.$$

From Theorem 1 it follows that for every graph G, the coefficient $c^\star_G$ *does not vanish* only if either $G\in\operatorname{mon}(\mathrm{BPM}^\star_\mathrm{n})$ or there exists some $H\in\operatorname{mon}(\mathrm{BPM}^\star_\mathrm{n})$ such that $G\in B(H)$. However, each of the aforementioned balls is relatively small (in fact, can be bounded by $|B(G)|\le 2^n\cdot n!$), thus by the union bound:

$$|\operatorname{mon}(\mathrm{UBPM}^\star_\mathrm{n})|\le|\operatorname{mon}(\mathrm{BPM}^\star_\mathrm{n})|\,(1+2^n\cdot n!)=2^{\Theta(n\log n)}$$

where the last equality follows from the bound $\log_2|\operatorname{mon}(\mathrm{BPM}^\star_\mathrm{n})|\le 2n\log_2 n+\mathcal{O}(n)$, obtained in [3]. This concludes the proof of the upper bound. The lower bound now follows directly from Theorem 1, as it suffices to observe that the coefficient of the complete bipartite graph is $\pm\operatorname{per}(K_{n,n})=\pm(n!)$. ◀

5 The Communication Rank of Unique Bipartite Matching

5.1 Rank and Polynomial Representation

The log-rank of a Boolean function is very closely related to its representation as a multilinear polynomial. This relationship is made very evident in the case of certain "lifted" functions: given a Boolean function $f : \{0,1\}^n \to \{0,1\}$, one can define the following pair of functions $f_\wedge, f_\oplus : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, where

$$\forall x, y \in \{0,1\}^n : f_\wedge(x,y) = f(x \wedge y), \text{ and } f_\oplus(x,y) = f(x \oplus y).$$

It is well known [18, 4] that the rank of the communication matrices $M_{f_\wedge}$ and $M_{f_\oplus}$ is *exactly characterized* by the *sparsity* (i.e., number of monomials) of the polynomials representing f in the $\{0,1\}$-basis and the $\{\pm 1\}$-basis (the Fourier basis), respectively. In other words,

$$\operatorname{rank}(M_{f_\wedge}) = \#\{\text{monomials in } \{0,1\}\text{-polynomial representing } f\}$$
$$\operatorname{rank}(M_{f_\oplus}) = \#\{\text{monomials in } \{-1,1\}\text{-polynomial representing } f\}.$$

The polynomial representation of a Boolean function f over the $\{0,1\}$-basis, or that of its dual $f^\star$, can also be used to derive communication rank upper bounds for non-lifted functions. The following lemma gives such a bound for the communication task of f, under *any* input partition.

▶ **Lemma 16.** *Let $f : \{0,1\}^n \to \{0,1\}$. Then, for every partition $S \sqcup \bar{S} = [n]$ we have:*

$$\operatorname{rank}(M_{f^{S \sqcup \bar{S}}}) \leq \min\left\{ |\operatorname{mon}(f)|, |\operatorname{mon}(f^\star)| + 1 \right\}.$$

Proof. Let $S \sqcup \bar{S} = [n]$ be some input partition, and let M and M' be the communication matrices of f and $f^\star$ under this partition, respectively. By definition of Boolean duality, we have $M = J - M_\pi M' M_\sigma$ where $J = \mathbb{1} \otimes \mathbb{1}$ is the all-ones matrix, and M_π, M_σ are the permutation matrices for

$$\forall x \subseteq S : \ \pi(x) = S \setminus x, \quad \forall y \subseteq \bar{S} : \ \sigma(y) = \bar{S} \setminus y$$

therefore $|\operatorname{rank}(M) - \operatorname{rank}(M')| \leq 1$, and it suffices to bound the rank of M'. However, we now observe that the polynomial representing f naturally induces a $|\operatorname{mon}(f)|$-rank decomposition of M (and likewise $f^\star$ for M'), as per [26], by considering the following sum of rank-1 matrices:

$$\forall T \in \operatorname{mon}(f), \text{ add the rank-1 matrix } a_T \cdot (\mathbb{1}_X \otimes \mathbb{1}_Y)$$

where a_T is the coefficient of T in f, and

$$X = \left\{ x : (T \cap S) \subseteq x \subseteq S \right\}, \quad Y = \left\{ y : \left(T \cap \bar{S}\right) \subseteq y \subseteq \bar{S} \right\}. \qquad \blacktriangleleft$$

5.2 The Rank of Unique Bipartite Matching

The log-rank of the unique bipartite matching function, ranging over all input partitions, is exactly characterized in the following Theorem.

▶ **Theorem 17.** *The log-rank of unique bipartite perfect matching is*

$$\max_{E \sqcup \bar{E} = E(K_{n,n})} \log \operatorname{rank}(M_{\mathrm{UBPM}_n^{E \sqcup \bar{E}}}) = \Theta(n \log n)$$

where $\mathrm{UBPM}_n^{E \sqcup \bar{E}}$ is the two-party function whose input is partitioned according to $E \sqcup \bar{E}$.

Proof. To obtain the lower bound, we must first fix a particular input partition. Assume without loss of generality that $n = 2m$ and let us partition the left and right vertices into two sets, $L = A \sqcup B$, $R = C \sqcup D$, where $A = \{a_1, \ldots, a_m\}$, $B = \{b_1, \ldots, b_m\}$, $C = \{c_1, \ldots, c_m\}$ and $D = \{d_1, \ldots, d_m\}$. Hereafter we consider the input partition wherein Alice receives all the edges incident to the left vertices A and Bob receives all the edges incident to the left vertices B. To prove our lower bound, we shall construct a *fooling set* (Definition 7). Let us introduce some notation: for every permutation $\pi \in S_m$ and two sets $X, Y \in \{A, B, C, D\}$, the notation $\pi(X, Y)$ refers to the matching from X to Y using the permutation π. Formally,

$$\forall X, Y \in \{A, B, C, D\} : \forall \pi \in S_m : \ \pi(X, Y) \stackrel{\text{def}}{=} \Big\{\{x_i, y_{\pi(i)}\} : i \in [m]\}\Big\}.$$

Under this notation, we claim that

$$S = \Big\{ \big(\mathrm{id}(A, C) \ \sqcup \ \pi(A, D), \ \ \mathrm{id}(B, D) \ \sqcup \ \big\{\{b_{\pi(i)}, c_j\} : 1 \leq i < j \leq m\big\}\big) \ : \ \pi \in S_m \Big\}$$

is a fooling set for $\mathrm{UBPM}_n^{K_{A,R} \sqcup K_{B,R}}$, where $\mathrm{id} \in S_m$ is the identity element.

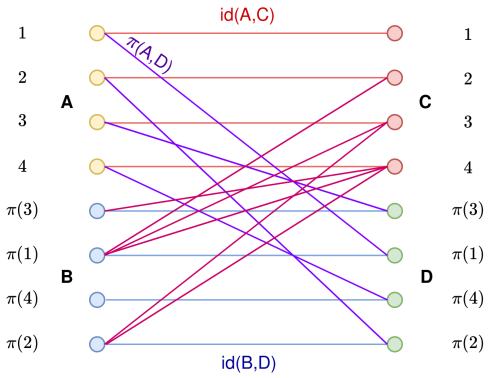

Figure 1 A graph G in the fooling set S, for $m = 4$ and $\pi = (2413)$.

$\{x \sqcup y : (x, y) \in S\} \subseteq \mathrm{UBPM}_n^{-1}(1)$: Let $\pi \in S_m$ and consider $G \subseteq K_{n,n}$ where:

$$E(G) = \mathrm{id}(A, C) \ \sqcup \ \pi(A, D) \ \sqcup \ \mathrm{id}(B, D) \ \sqcup \ \big\{\{b_{\pi(i)}, c_j\}\big\}_{i<j}.$$

Clearly G has the identity perfect matching, whereby A is matched to C and B to D. Let us denote this matching by M. To show that M is *unique*, it suffices to show that there exists no M-alternating cycle in G. By construction, the vertices in any such cycle must alternate between $C - A - D - B$ (since the only edges joining $A \leftrightarrow C$ and $B \leftrightarrow D$ are those in the matching M). Thus, for any $i \in [m]$, an M-alternating path starting with c_i must be of the form:

$$c_i \sim a_i \sim d_{\pi(i)} \sim b_{\pi(i)} \sim c_j \sim \ldots$$

where $j > i$. However, observe that $b_{\pi(m)}$ is not adjacent to any vertex in C, so any such path will eventually (after at most m passes through B) terminate at $b_{\pi(m)}$, without looping back to c_i. Therefore there exists no M-alternating cycle, and M is indeed unique.

$\forall (x_1, y_1), (x_2, y_2) \in S : \ (x_1 \sqcup y_2) \in \mathrm{UBPM_n}^{-1}(0)$: Let $\pi, \sigma \in S_m$ where $\pi \neq \sigma$, and let G be the graph:

$$E(G) = \mathrm{id}(A, C) \ \sqcup \ \pi(A, D) \ \sqcup \ \mathrm{id}(B, D) \ \sqcup \ \big\{\{b_{\sigma(i)}, c_j\}\big\}_{i<j}.$$

Once again, clearly G has the identity matching M, whereby A is matched to C and B to D. To show that M is not unique, it suffices to exhibit an alternating cycle. Recall that $\sigma \neq \pi$ and therefore $\sigma^{-1} \circ \pi \neq \mathrm{id}$, and in particular, there exists some $i \in [m]$ such that $\sigma^{-1}(\pi(i)) < i$. By construction, the following M-alternating cycle is present in G:

$$c_i \sim a_i \sim d_{\pi(i)} \sim b_{\pi(i)} = b_{\sigma(\sigma^{-1}(\pi(i)))} \sim c_i.$$

Therefore, S is a fooling set for $\mathrm{UBPM_n}$ under the aforementioned input partition. To conclude the lower bound, we recall the following Theorem, due to Dietzfelbinger, Hromkovič and Schnitger [8]:

▶ **Theorem 18** ([8]). *$\forall f : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}$ we have $\log_2 \mathrm{fs}(f) \leq 2\left(\log_2 \operatorname{rank} M_f + 1\right)$.*

Therefore, we have:

$$\log_2 \operatorname{rank}\left(M_{\mathrm{UBPM}_n^{K_{A,R} \sqcup K_{B,R}}}\right) \geq \tfrac{1}{2}\log_2 |S| - 1 = \tfrac{1}{4} n \log_2 n - \Theta(n)$$

concluding the lower bound. As for the upper bound, it follows directly from Lemma 16, and from the characterization of Theorem 2 (see Corollary 15). ◀

References

1. Scott Aaronson, Shalev Ben-David, Robin Kothari, Shravas Rao, and Avishay Tal. Degree vs. approximate degree and quantum implications of Huang's sensitivity theorem. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 2021.
2. Gal Beniamini. The approximate degree of bipartite perfect matching. *arXiv preprint*, 2020. `arXiv:2004.14318`.
3. Gal Beniamini and Noam Nisan. Bipartite perfect matching as a real polynomial. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 2021.
4. Anna Bernasconi and Bruno Codenotti. Spectral analysis of boolean functions as a graph eigenvalue problem. *IEEE transactions on computers*, 48(3):345–351, 1999.
5. Louis J Billera and Aravamuthan Sarangarajan. The combinatorics of permutation polytopes. In *Formal power series and algebraic combinatorics*, volume 24, pages 1–23, 1994.
6. Harry Buhrman and Ronald De Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288(1):21–43, 2002.
7. Mark Bun and Justin Thaler. Guest column: Approximate degree in classical and quantum computing. *ACM SIGACT News*, 51(4):48–72, 2021.
8. Martin Dietzfelbinger, Juraj Hromkovič, and Georg Schnitger. A comparison of two lower-bound methods for communication complexity. *Theoretical Computer Science*, 168(1):39–51, 1996.
9. Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17:449–467, 1965.
10. Stephen Fenner, Rohit Gurjar, and Thomas Thierauf. Bipartite perfect matching is in Quasi-NC. *SIAM Journal on Computing*, 50(3):STOC16–218, 2019.
11. Harold N Gabow, Haim Kaplan, and Robert E Tarjan. Unique maximum matching algorithms. *Journal of Algorithms*, 40(2):159–183, 2001.
12. Martin Charles Golumbic, Tirza Hirst, and Moshe Lewenstein. Uniquely restricted matchings. *Algorithmica*, 31(2):139–154, 2001.
13. Gábor Hetyei. Rectangular configurations which can be covered by 2×1 rectangles. *Pécsi Tan. Foisk. Közl*, 8:351–367, 1964.
14. Thanh Minh Hoang, Meena Mahajan, and Thomas Thierauf. On the bipartite unique perfect matching problem. In *International Colloquium on Automata, Languages, and Programming*, pages 453–464. Springer, 2006.

15 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.

16 Hao Huang. Induced subgraphs of hypercubes and a proof of the sensitivity conjecture. *Annals of Mathematics*, 190(3):949–955, 2019.

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

18 Alexander Knop, Shachar Lovett, Sam McGuire, and Weiqiang Yuan. Log-rank and lifting for AND-functions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 197–208, 2021.

19 Dexter Kozen, Umesh V Vazirani, and Vijay V Vazirani. NC algorithms for comparability graphs, interval graphs, and testing for unique perfect matching. In *International Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 496–503. Springer, 1985.

20 Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, USA, 1996.

21 László Lovász. On determinants, matchings, and random algorithms. In *FCT*, volume 79, pages 565–574, 1979.

22 László Lovász and Michael Saks. Lattices, Möbius functions and communications complexity. In *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, pages 81–90. IEEE Computer Society, 1988.

23 Laszlo Lovasz and Neal E Young. Lecture notes on evasiveness of graph properties. *arXiv preprint cs/0205031*, 2002.

24 Kurt Mehlhorn and Erik M Schmidt. Las vegas is better than determinism in VLSI and distributed computing. In *Proceedings of the fourteenth annual ACM symposium on Theory of computing*, pages 330–337, 1982.

25 Noam Nisan. The demand query model for bipartite matching. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 592–599. SIAM, 2021.

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

27 Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.

28 M.D. Plummer and L. Lovász. *Matching Theory*. North-Holland Mathematics Studies. Elsevier Science, 1986.

29 Alexander A Sherstov. Algorithmic polynomials. *SIAM Journal on Computing*, 49(6):1173–1231, 2020.

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

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

A The Approximate Degree of $\mathrm{UBPM_n}$

The ε-approximate degree $\widetilde{\deg}_\epsilon(f)$, of a Boolean function $f : \{0,1\}^n \to \{0,1\}$ is the *least* degree of a real multilinear polynomial *approximating* f pointwise over $\{0,1\}^n$, with error at most ε. Formally,

▶ **Definition 19.** *Let $f : \{0,1\}^n \to \{0,1\}$ and let $0 < \varepsilon < \frac{1}{2}$. The ε-approximate degree of f, $\widetilde{\deg}_\epsilon(f)$, is the least degree of a real multilinear polynomial $p \in \mathbb{R}[x_1, \ldots, x_n]$ such that:*

$$\forall x \in \{0,1\}^n : \ |f(x) - p(x)| \leq \varepsilon.$$

If $\varepsilon = 1/3$, then we omit the subscript in the above notation, and instead write $\widetilde{\deg}(f)$.

Approximate degree is a well-studied complexity measure. For a comprehensive survey on the topic, we refer the reader to [7]. With regards to Theorem 2, we make the following observation: every Boolean function whose polynomial representation, or that of its dual, have low ℓ_1-norm – can be efficiently *approximated* in the ℓ_∞-norm by a low-degree polynomial. Firstly, it is not hard to see that for any Boolean function $f : \{0,1\}^n \to \{0,1\}$ and any $\varepsilon > 0$, the ε-approximate degree of f is identical to that of its dual $f^\star$. This follows since $f^\star$ can be obtained through an *affine transformation* of f, which cannot increase the degree, and the same transformation can similarly be applied to any approximating polynomial of f (and the converse follows since $(f^\star)^\star = f$). The second component of the approximation scheme is the following lemma.

▶ **Lemma 20** ([2], similar to [29]). *Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function, and let $p \in \mathbb{R}[x_1, \ldots, x_n]$ be its representing polynomial, where $\|p\|_1 \in [3, 2^n]$. Then:*

$$\forall \|p\|_1^{-1} \leq \varepsilon \leq \frac{1}{3} : \quad \widetilde{\deg}_\epsilon(f) = \mathcal{O}\left(\sqrt{n \log \|p\|_1}\right).$$

The proof of Lemma 20 follows from the following simple approximation scheme: replace every *monomial* (of sufficiently large degree) with a polynomial that approximates it pointwise, to some sufficiently small error (depending only on the ℓ_1-norm of the representing polynomial). The full details of this scheme appeared previously in [2, 29]. Combining Lemma 20 with the ℓ_1-bound of Corollary 15, we obtain:

▶ **Corollary 21.** *For any $n > 1$, and $2^{-n \log n} \leq \varepsilon \leq \frac{1}{3}$, we have:*

$$\widetilde{\deg}_\epsilon(\mathrm{UBPM_n}) = \mathcal{O}(n^{3/2}\sqrt{\log n}).$$

B Families of Matching Functions having Low Dual ℓ_1-Norm

The main algorithmic result in this paper is the low ℓ_1-norm of the dual function of $\mathrm{UBPM_n}$, from which we deduce *upper bounds*, for instance on the communication rank and the approximate degree. In [2], a similar bound had been obtained for the dual of the perfect matching function, $\mathrm{BPM_n}$. These norm bounds and their corollaries extend to a wide range of *matching-related functions*, some of which are detailed below.

Functions Obtained by Restrictions

Consider any two Boolean functions f and g, such that g is obtained by a *restriction* of f (i.e., by fixing some of the inputs bits of f). As restrictions cannot increase the norm, it clearly holds that $\|g\|_1 \leq \|f\|_1$ and $\|g^\star\|_1 \leq \|f^\star\|_1$. Several intrinsically interesting matching-functions can be cast in this way. One notable example is the bipartite k-matching function, which is the indicator over all graphs $G \subseteq K_{n,n}$ containing a matching of size k.

$$\mathrm{BM_{n,k}}(x_{1,1}, \ldots, x_{n,n}) = \begin{cases} 1 & \{(i,j) : x_{i,j} = 1\} \subseteq K_{n,n} \text{ has a } k\text{-matching} \\ 0 & \text{otherwise.} \end{cases}$$

This function is obtained by a restriction of BPM_{2n-k}, as follows. Label the vertices of $K_{2n-k,2n-k}$ by

$$L = A \sqcup V, \text{ where } A = \{a_1, \ldots, a_n\}, V = \{v_1, \ldots, v_{n-k}\}$$
$$R = B \sqcup U, \text{ where } B = \{b_1, \ldots, b_n\}, U = \{u_1, \ldots, u_{n-k}\}.$$

Given any input $G \subseteq K_{n,n}$ to $\mathrm{BM_{n,k}}$, the edges of G are encoded via the edges joining A and B, and moreover we fix two additional bicliques $K_{A,U}$, $K_{V,B}$. The resulting graph contains a bipartite perfect matching if and only if G contains a k-matching, and thus

▶ **Corollary.** *For every* $0 < k \leq n$*, we have* $\log \left\|\mathrm{BM}^{\star}_{\mathrm{n,k}}\right\|_1 = \mathcal{O}(n \log n)$.

This norm bound is tight whenever $k = \alpha n$, for any *constant* $0 < \alpha < 1$, as are (up to log-factors) the bounds on the approximate degree and on the log-rank.

▶ **Corollary.** *Let* $\alpha \in (0,1)$ *be a constant. Then for every* $n > 1$ *and* $2^{-n \log n} \leq \varepsilon \leq \frac{1}{3}$*, we have:*

$$\log \left\|\mathrm{BM}^{\star}_{n,\alpha n}\right\|_1 = \Theta(n \log n), \quad \widetilde{\deg}_{\epsilon}\left(\mathrm{BM}_{n,\alpha n}\right) = \widetilde{\Theta}(n^{3/2}), \text{ and } \log \operatorname{rank}\left(\mathrm{BM}_{n,\alpha n}\right) = \widetilde{\Theta}(n).$$

The aforementioned approximate degree lower bound follows using the method of *Spectral Sensitivity* – a complexity measure due to Aaronson, Ben-David, Kothari, Rao and Tal [1], based on Huang's proof of the sensitivity conjecture [16]. [1] proved that the approximate degree of any total function f is bounded below by the spectral radius of its *sensitivity graph* (i.e., the f-cut of the hypercube). As this graph is bipartite, its spectrum is symmetric, and it therefore suffices (by Cauchy interlacing) to obtain a lower bound on the spectral radius of any vertex induced subgraph of the sensitivity graph [2].

For $\mathrm{BM_{n,k}}$ this construction is straightforward – consider the induced graph whose left vertices are all $(k-1)$-matchings, and right vertices are all k-matchings. This produces a biregular subgraph of the sensitivity graphs of $\mathrm{BM_{n,k}}$, with left degrees $(n-k+1)^2$ and right degrees k. As it is well known that the spectral radius of a biregular graph is $\sqrt{d_L d_R}$ (where d_L and d_R are the left and right degrees, respectively), this concludes the bound on the spectral sensitivity of $\mathrm{BM_{n,k}}$, and by extension, its approximate degree[4]. This lower bound on $\widetilde{\deg}\left(\mathrm{BM_{n,k}}\right)$ now implies the ℓ_1-norm lower bound, through Lemma 20.

As for the log-rank lower bound, it follows by a simple fooling set argument, under the same input partition used in Theorem 17. Let $L = A \sqcup B$ be the left vertices corresponding to the input partition, where $|A| = |B| = n/2$, and let A' and B' be the first $k/2$ vertices of A and B, respectively. Let C be the first $k = \alpha n$ right vertices. Then, under the notation of Theorem 17,

$$S = \left\{ \left(\mathrm{id}(A', S), \mathrm{id}(B', \bar{S})\right) \ : \ S \subseteq C, \ \bar{S} = C \setminus S, \ |S| = |\bar{S}| = \tfrac{k}{2} \right\}$$

is a fooling set for $\mathrm{BM}_{n,\alpha n}$, where the indices of S and $\bar{S}$ correspond to a *fixed* ordering on C. Any pair (x, y) contains a k-matching, but for any mismatching pair belonging to sets $S_1 \neq S_2 \subseteq C$, we have that $S_1 \cap S_2 \neq \emptyset$ and thus the maximum matching is of size $|S_1 \cup S_2| < k$. By construction, this fooling set is of size

$$\log_2 |S| = \log_2 \binom{k}{k/2} = k - o(1)$$

and the log-rank bound now follows from Theorem 18.[5]

[4] We remark that the same construction also trivially extends to $\mathrm{UBM_{n,k}}$; the *unique* k-matching function.

[5] For the *unique* bipartite k-matching function $\mathrm{UBM_{n,k}}$ one can obtain a slightly stronger log-rank bound by repeating the construction of Theorem 17 with k-matchings rather than perfect matchings, and by adding $n-k$ isolated vertices. This yields a log-rank bound of $\log_2 (k/2!) = \Theta(k \log k)$.

Formulas over Low-Norm Functions

Given any two nontrivial Boolean functions f and g, the norms of their conjunction, disjunction, and negation are at-most multiplicative in their respective norms, and the same holds for their duals. Therefore, the dual of any short De Morgan formula whose atoms are Boolean functions of low dual ℓ_1-norm, will similarly inherit the low-norm property. Several matching functions can be represented in this way, and thus have low dual norm. For example

$$\mathrm{MaxMatch}_{\mathrm{n,k}}(x_{1,1},\ldots,x_{n,n}) = \begin{cases} 1 & \text{The maximum matching of } \left\{(i,j) : x_{i,j} = 1\right\} \text{ is of size } k \\ 0 & \text{otherwise} \end{cases}$$

can be constructed as $\mathrm{MaxMatch}_{\mathrm{n,k}} = \mathrm{BM}_{\mathrm{n,k}} \wedge \neg\, \mathrm{BM}_{n,k+1}$.