The Combinatorial Nullstellensatz and DFT on Perfect Matchings in Bipartite Graphs

Timothy M. Brauch
Department of Mathematics
University of Louisville
Louisville, KY 40292 USA
timothy.brauch@louisville.edu

André E. Kézdy
Department of Mathematics
University of Louisville
Louisville, KY 40292 USA

Hunter S. Snevily
Department of Mathematics
University of Idaho
Moscow, ID 83844 USA

February 8, 2009

Abstract

The paper begins with a simple circular lock problem that shows how the Combinatorial Nullstellensatz relates to the discrete Fourier Transform. Specifically, the lock shows a relationship between detecting perfect matchings in bipartite graphs using the Combinatorial Nullstellensatz and detecting a maximum rank independent set in the intersection of two matroids in the Fourier transform of a specially chosen function. Finally, an application of the uncertainity principle computes a lower bound for the product of perfect matchings and the number of independent sets.

1 Introduction

Imagine an $n \times n$ *circular lock* consisting of n equal-sized wheels placed one on top of the other where each wheel has n cells. All of the cells are the same size and are filled with a complex number. Each wheel rotates independently, both clock-wise and counter-clockwise, but only in discrete intervals corresponding to the cell sizes so that, after rotations are complete, cells align forming columns. A *setting* of the lock is such a rotation of the wheels. Because a setting of the lock means that the cells align, each setting determines (up to a rotation of all wheels

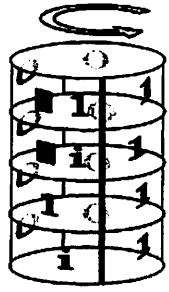

$$\begin{bmatrix} -1 & 1 & 0 & 0 \\ -i & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ i & 1 & 0 & 0 \end{bmatrix}$$

Figure 1: A setting on a circular lock and the corresponding matrix.

by the same number of cells) an $n \times n$ matrix whose ijth entry is simply the entry of the jth cell of wheel i. A setting of a 4×4 circular lock is shown in Figure 1.

A circular lock is unlocked or *opens* if its wheels are placed into a setting in which the corresponding matrix has nonzero determinant; otherwise, the lock remains *closed*. This raises a few questions. Can a given lock be opened at all? If so, which setting opens it? How many settings are there that open it? Is there a small set of "master" settings to open all locks?

As we shall see, the first two "circular-lock" questions arise naturally from the consideration of a polynomial designed to detect perfect matchings in bipartite graphs and its Fourier transform which detects maximum cardinality independent sets in the intersection of two specific types of matroids. The circular locks this paper investigates are the ones whose rows are coefficients of polynomials with zeros that are all nth roots of unity. The third question is discussed in the paper in the last section while the fourth question is an area of future research.

The motivation to study perfect matchings in bipartite graphs in this setting arose from our frustration in applying the Combinatorial Nullstellensatz to a host of famous open problems. The Combinatorial Nullstellensatz is a tool that detects the existence of a combinatorial object by showing that a certain polynomial does not vanish over some domain.

The most applicable version of the Combinatorial Nullstellensatz is

Theorem 1 (Combinatorial Nullstellensatz [1]) *Let F be an arbitrary field, and let f be a polynomial in $F[x_1, \ldots, x_n]$. Suppose the degree of f is $\sum_{i=1}^{n} t_i$, where each t_i is a nonnegative integer. If the coefficient of $\prod_{i=1}^{n} x_i^{t_i}$ in f is nonzero, then for any subsets $S_1, S_2, \ldots, S_n$ of F satisfying $|S_i| > t_i$, there are elements $s_1 \in S_1, s_2 \in S_2, \ldots, s_n \in S_n$ such that*

$$f(s_1, s_2, \ldots, s_n) \neq 0.$$

The Combinatorial Nullstellensatz is closely related to Fourier transforms over a finite group and this relationship will be explored later in this paper. The Com-

binatorial Nullstellensatz has been employed successfully in a variety of circumstances, but there is still no clear understanding of which circumstances are favorable to its application despite the many problems that are apparently prime candidates. One of the main purposes of this paper is to show how it can be applied to the problem of detecting perfect matchings in bipartite graphs. Other problems for which it seems aptly suited include: the problem of showing that every tree has a ρ-valuation (see [5]), showing that every odd order latin square has a latin transversal [6], and proving the existence of a hamiltonian cycle in middle levels of the boolean lattice, just to name a few of the highly symmetric, famous and still open problems.

In each of these problems it is straight forward to construct polynomials that vanish completely on some appropriate domain if and only if the desired combinatorial object does not exist. The main source of our frustration is the realization that the polynomials in question are presented in compact, factored form; determining whether a nonzero coefficient appears in its expansion (modulo an appropriate ideal) is a formidable problem (in general, this problem is NP-hard). Most successful applications of the nullstellensatz technique so far, when applied to problems with more than one instance of each size, have been to problems with the special property that all instances of a given size determine a collection of polynomials that have a common monomial with a nonzero coefficient; thus proving the monomial is nonzero for one canonical instance shows it is nonzero for others. Many natural problem formulations do not share this property. The natural formulation of the ρ-valuations-for-trees problem, for example, does not have this property. Similarly, the natural formulation of the perfect-matching-in-a-bipartite graph problem also does not, as we shall see. Because this latter problem is easy from a complexity point of view (which is why we chose it for investigation), one would expect a polynomial-time algorithm to find a nonzero coefficient in the expansion of the corresponding encoding polynomial, if such a coefficient exists. The matroid-intersection algorithm suffices for this purpose, as we shall see in the second section.

We hope that further investigation will provide a sharpened form of the Combinatorial Nullstellensatz, perhaps incorporating elements of the matroid-intersection algorithm. It seems very likely a nice formulation along these lines awaits discovery. This paper demonstrates that such a formulation applies in the perfect-matching-in-a-bipartite graph problem. Formulating and solving this problem in the nullsellensatz fashion has shed some light on the relation between the number of perfect matchings and the number of maximum independent sets in the intersection of certain matroids via the uncertainty principle, as formulated through the Fourier transform on a finite group. We explain this in the last section.

2 Circular Locks from Bipartite Graphs

In this section we formalize our "Combinatorial Nullstellensatz" approach to detecting perfect matchings in bipartite graphs. It should be noted that other approaches are possible (for example using other polynomials or domains), but the approach we take leads naturally to the discrete Fourier transform and so shares many of its desirable qualities.

Let $\mathbb{C}$ denote the field of complex numbers and $\omega = e^{2\pi i/n}$ where $i = \sqrt{-1}$. For a positive integer n, let $\Omega_n = \{\omega^0, \ldots, \omega^{n-1}\}$, be the set of nth roots of unity.

Consider a bipartite graph G with vertex set $A \cup B$, where $A = \{0, \ldots, n-1\}$ and $B = \Omega_n$, and edge set $E \subseteq \{\{a, b\} : a \in A, b \in B\}$. Recall that a *matching* in a graph is a collection of disjoint edges. A matching is *perfect* if every vertex of the graph is contained in an edge of the matching. A classical problem in graph theory is to characterize bipartite graphs that have a perfect matching. Numerous theorems (e.g. Hall's theorem) and algorithms (e.g. the alternating path algorithm) have been developed to solve this problem.

For $i = 0, \ldots, n-1$, introduce a variable x_i and a univariate polynomial

$$g_i(x_i) = \prod_{\{i,\omega^j\} \notin E} (x_i - \omega^j),$$

where it is understood that the product is 1 if vertex i is adjacent to all vertices in B. The polynomial G_i has degree at most $n-1$ as long as vertex i is not an isolated vertex; however, as an isolated vertex does not admit a perfect matching the results contained in this paper will not apply. We will assume there are no isolated vertices for the rest of this paper.

Recall the Vandermonde identity,

$$V(x) = V(x_0, \ldots, x_{n-1}) = \prod_{0 \le i < j < n} (x_j - x_i) = \sum_{\pi \in \mathcal{S}_n} \operatorname{sign}(\pi) \prod_{i=0}^{n-1} x_i^{\pi(i)},$$

where $\mathcal{S}_n$ is the set of permutations of $0, \ldots, n-1$.

There exists a perfect matching in G if and only if the polynomial

$$f_G(x_0, \ldots, x_{n-1}) = \prod_{0 \le i < j < n} (x_j - x_i) \prod_{i=0}^{n-1} g_i(x_i)$$

is nonzero for some input from Ω_n^n. We shall use $f(x)$ as an abbreviation for $f_G(x_0, \ldots, x_{n-1})$ or for clarity about which graph is being discussed, $f_G(x)$. Following common usage, for any $\alpha \in \mathbb{Z}_n^n$, we shall use x^α as an abbreviation for $\prod_{i=0}^{n-1} x_i^{\alpha_i}$.

Expand the polynomials g_i into sums,

$$g_i(x_i) = \prod_{\{i,\omega^j\} \notin E} (x_i - \omega^j) = \sum_{j=0}^{n-1} \ell_{ij} x_i^j,$$

$$\begin{bmatrix} -1 & 0 & 1 & 0 \\ i & -1-i & 1 & 0 \\ -i & -i+1 & 1 & 0 \\ i & 1+i & 1 & 0 \end{bmatrix}$$

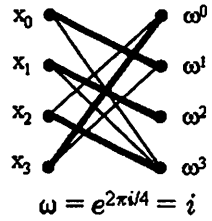

Figure 2: A circular lock derived from its bipartite graph.

and let $L_G = [\ell_{ij}]$ be the $n \times n$ matrix of coefficients. We view L_G as the circular lock derived from G and correspondingly G_L is the graph related to the lock L. Given any $\alpha \in \mathbb{Z}_n^n$, let $L[\alpha]$ denote the matrix obtained from L by rotating, for each $i = 1, \ldots, n$, row i to the left α_i units with wrap around.

Observe that G_L does not depend on the setting of the lock L_G because, for all $i, j \in \{0, \ldots, n-1\}$, the value ω^j is a root of $g_i(x_i)$ if and only if ω^j is a root of $x_i^k g_i(x_i)$ modulo $x_i^n - 1$, for all integers k. Though many locks determine the same bipartite graph by rotating the rows, each bipartite graph G on the vertices $\{0, \ldots, n-1\} \cup \Omega_n$ can be associated with a canonical circular lock. This lock is the one in which, before rotation, the entries of the cells on wheel i correspond to the coefficients of the polynomial having precisely those elements of Ω_n as roots that coincide with the neighbors of vertex i in G (see Figure 2).

Define $\mathcal{I}_n$ to be the ideal in $\mathbb{C}[x_0, \ldots, x_{n-1}]$ consisting of those polynomials that vanish on all inputs from the nth roots of unity; that is $f \in \mathcal{I}_n$ if and only if $f(a) = 0$, for all $a \in \Omega_n^n$. In [6] Kézdy and Snevily show that $\mathcal{I}_n = \langle x_i^n - 1 \rangle_{i=0}^{n-1}$.

Theorem 2 *A circular lock L_G opens if and only if G_L has a perfect matching.*

Proof. Recall that there exists a perfect matching in $G = G_L$ if and only if the polynomial

$$f_G(x) = V(x) \prod_{i=0}^{n-1} g_i(x_i) \tag{1}$$

is nonzero for some input from Ω_n^n. Now consider the polynomial

$$g(x) = f_G(x) \text{ modulo } \mathcal{I}_n.$$

We first prove that

$$g(x) = \sum_{\alpha \in \mathbb{Z}_n^n} \det(L[\alpha]) x^\alpha \tag{2}$$

To prove (2), it suffices to prove that, for all $\alpha \in \mathbb{Z}_n^n$, the constant coefficient of $x^{-\alpha}g(x)$ modulo $\mathcal{I}_n$ is $\det(L[\alpha])$ (all exponents are reduced modulo n). Now the computation

$$x^{-\alpha}g(x) \equiv V(x)\left(\prod_{i=0}^{n-1} x_i^{-\alpha_i} g_i(x_i)\right) \text{ modulo } \mathcal{I}_n \tag{3}$$

shows that multiplying $g(x)$ by $x^{-\alpha}$ has the effect (modulo $\mathcal{I}_n$) of shifting, for all i, the coefficients of each $g_i(x_i)$ to the left by α_i units in the matrix L. To obtain the constant coefficient of $x^{-\alpha}g(x)$, observe that the Vandermonde polynomial expands into $\sum_{\pi \in S_n} \text{sign}(\pi) \prod_{i=0}^{n-1} x_i^{\pi(i)}$ so, in order to obtain a constant coefficient, a monomial must be chosen from each of the factors $x_i^{-\alpha_i} g_i(x_i)$ that appear in (3) in such a way that no two monomials have the same exponent; that is, we must select a transversal in the matrix $L[\alpha]$. This, along with the weighting of permutations by signs that appears in the expansion of the Vandermonde polynomial, means that the constant coefficient of $x^{-\alpha}g(x)$ modulo $\mathcal{I}_n$ is $\det(L[\alpha])$.

Because $g(x)$ is equivalent to $f(x)$ which has the form 1, it is clear that $g(\alpha) \neq 0$, for some $\alpha \in \Omega_n^n$, if and only if G_L has a perfect matching. On the other hand, form 2 of $g(x)$ shows that $g \not\equiv 0$ modulo $\mathcal{I}_n$, if and only if $\det(L[\alpha]) \neq 0$, for some $\alpha \in \mathbb{Z}_n^n$; the theorem follows. ◇

As an example of Theorem 2, the circular lock depicted in Figure 1 corresponds to the complete bipartite graph $K_{4,4}$ minus the edges of a perfect matching; so it opens. Indeed, the setting obtained by rotating the bottom two rows by two units to the right opens the lock. Many other settings open the lock too.

3 Matroids and Matroid Intersection

Theorem 2 gives an efficient means to test whether a given circular lock opens, but gives no efficient means to find which setting opens a circular lock that does open. Efficiently finding such a setting can be accomplished by employing matroid intersection.

For our purposes it suffices to mention two fundamental examples of matroids. First, if E is a set of vectors over a field F, and $\mathcal{I}$ contains the empty set together with those subsets of E that form linearly independent sets of vectors over F, then $(E, \mathcal{I})$ is a matroid; such a matroid is called a *vector matroid*. Second, if E is a finite set with a partition into subsets $E_1, \ldots, E_t$ and $\mathcal{I} = \{A \subseteq E : |A \cap E_i| \leq 1, \text{ for all } i = 1, \ldots, t\}$, then $(E, \mathcal{I})$ is a matroid; such a matroid is called a *partition matroid*.

The matroid intersection algorithm, as can be seen in [9], computes the maximum cardinality of an independent set common to two matroids. The classic book by Lawler [7] and the more recent book by Cook, Cunningham, Pulleyblank, and

Schrijver [2]) contain introductions to polynomial-time matroid intersection algorithms which we use here.

Theorem 3 *Let $G(A, B; E)$ be a bipartite graph and f_G its corresponding polynomial as defined in (1). The settings that open the circular lock L_G correspond to nonzero coefficients in the polynomial f_G and such a setting (if it exists) can be found in polynomial time via matroid intersection.*

Proof. Consider a circular lock $L = L_G$, and its corresponding polynomial $f_G(x)$ as in equation (1). As in the proof of Theorem 2 is suffices to show it is true for the equivalent equation $g(x)$ of the form given in equation (2). The coefficients of g correspond to determinants arising from rotations of the rows of L. Let $r_1, \ldots, r_n$ be the row vectors of L. Define E_j as the set of the n vectors obtained from r_j by cyclically permuting coordinates. Now define two matroids, M_1 and M_2 on the common ground set $E = \cup_{j=1}^{n} E_j$. The matroid M_1 is the vector matroid on E in which a set of elements is independent if and only if they are linearly independent (over $\mathbb{C}$). The matroid M_2 is the partition matroid on E in which a set of elements$S \subseteq E$ is independent if and only if $|S \cap E_j| \leq 1$, for all $j = 1, \ldots, n$. Now there is some $\alpha \in \mathbb{Z}_n^n$ such that $\det(L[\alpha]) \neq 0$ if and only if there exists a common independent set in M_1 and M_2 with cardinality n. Detecting the existence of such an independent set (and constructing such a set, if it exists) can be accomplished in polynomial time by Edmond's Matroid Intersection Algorithm. ◇

Theorem 3 gives an efficient means of finding a nonzero coefficient in the expansion of $g(x)$. Detecting perfect matchings in bipartite graphs via matroid intersection is not particularly novel. However, the interesting aspect of Theorem 3 is that it shows the connection between the coefficients of g that correspond to cardinality n sets in the intersection of M_1 and M_2, and the valuations of g over Ω_n^n that correspond to perfect matchings in G_L.

4 The Fourier Transform on Ω_n^n

This section provides background material and formulas that will be important for the result in the next section. In particular, equation (5) is necessary for a short proof of Theorem 4. In [11], Terras provides much more background on this subject.

The Fourier matrix of order n is the $n \times n$ matrix

$$F_n = \frac{1}{\sqrt{n}} \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \xi & \xi^2 & \cdots & \xi^{n-1} \\ 1 & \xi^2 & \xi^4 & \cdots & \xi^{2(n-1)} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & \xi^{n-1} & \xi^{2(n-1)} & \cdots & \xi^{(n-1)(n-1)} \end{pmatrix}$$

where $\xi = e^{-2\pi i/n}$ and $i = \sqrt{-1}$. Observe that F_n is symmetric and invertible. The inverse of F_n is its conjugate transpose, denoted F_n^*, and can be obtained from F_n by replacing ξ with $\omega = e^{2\pi i/n}$.

The Vandermonde matrix is defined as

$$V(z_0, \ldots, z_{n-1}) = \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ z_0 & z_1 & z_2 & \cdots & z_{n-1} \\ z_0^2 & z_1^2 & z_2^2 & \cdots & z_{n-1}^2 \\ \vdots & \vdots & \vdots & & \vdots \\ z_0^{n-1} & z_1^{n-1} & z_2^{n-1} & \cdots & z_{n-1}^{n-1} \end{pmatrix}.$$

The Fourier matrix F_n can be written as $F_n = n^{-1/2}V(\xi^0, \xi^1, \ldots, \xi^{n-1})$ by making the appropriate substitutions in the Vandermonde matrix.

It is well known that if $p(z) = \sum_{j=0}^{n-1} a_j z^j$ is a polynomial of degree at most $n-1$, then $p(z)$ is determined uniquely by its value at n distinct points. In particular, if these points are the nth roots of unity, then

$$n^{1/2}F_n^* \begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_{n-1} \end{pmatrix} = \begin{pmatrix} p(\omega^0) \\ p(\omega^1) \\ \vdots \\ p(\omega^{n-1}) \end{pmatrix}$$

from which it follows that

$$\begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_{n-1} \end{pmatrix} = n^{-1/2}F_n \begin{pmatrix} p(\omega^0) \\ p(\omega^1) \\ \vdots \\ p(\omega^{n-1}) \end{pmatrix}. \tag{4}$$

Let $G = V(A, B; E)$ be a bipartite graph, where $A = \{0, 1, \ldots, n-1\}$, $B = \Omega_n$. Recall from Section 2 that

$$f_G(x_0, \ldots, x_{n-1}) = \prod_{0 \le i < j < n} (x_j - x_i) \prod_{i=0}^{n-1} g_i(x_i)$$

is nonzero for some input from Ω_n^n if and only if there is a perfect matching in G.

Now the Fourier transform of $f = f_G : \Omega_n^n \to \mathbb{C}$ is the function

$$\hat{f}(\chi) = \sum_{\alpha \in \Omega_n^n} f(\alpha)\overline{\chi(\alpha)},$$

where χ is an element of the dual of Ω_n^n. Because Ω_n^n is abelian, the dual of Ω_n^n is isomorphic to itself. If $\alpha = (\alpha_1, \ldots, \alpha_n)$ and $\chi = (\chi_1, \ldots, \chi_n)$ where

$\chi_i = \omega^{r(i)}$, then

$$\overline{\chi(\alpha)} = \prod_{i=1}^{n} \alpha_i^{-r(i)}.$$

Simplifying the Fourier transform reveals a result that will prove useful later in the paper. Define P_π as

$$P_\pi = [p_{i,j}] = \begin{cases} 1 & \text{if } \pi(i) = j, \\ 0 & \text{otherwise.} \end{cases}$$

Then,

$$\begin{aligned}
\hat{f}(\chi) &= \sum_{\alpha \in \Omega_n^n} f(\alpha)\overline{\chi(\alpha)} \\
&= \sum_{\alpha \in \Omega_n^n} \det(V(\alpha_1, \ldots, \alpha_n)) \left(\prod_{i=1}^{n} g_i(\alpha_i) \right) \left(\prod_{i=1}^{n} \alpha_i^{-r(i)} \right) \\
&= \sum_{\pi \in \mathcal{S}_n} \det(V(\omega^{\pi_1}, \ldots, \omega^{\pi_n})) \left(\prod_{i=1}^{n} (\omega^{\pi_i})^{-r(i)} g_i(\omega^{\pi_i}) \right) \\
&= \sum_{\pi \in \mathcal{S}_n} \det(V(\xi^0, \ldots, \xi^{n-1}) \cdot P_\pi) \left(\prod_{i=1}^{n} (\omega^{\pi_i})^{-r(i)} g_i(\omega^{\pi_i}) \right) \\
&= \sum_{\pi \in \mathcal{S}_n} \det(n^{1/2} F_n) \det(P_\pi) \left(\prod_{i=1}^{n} (\omega^{\pi_i})^{-r(i)} g_i(\omega^{\pi_i}) \right) \\
&= n^{n/2} \det(F_n) \sum_{\pi \in \mathcal{S}_n} \det(P_\pi) \left(\prod_{i=1}^{n} (\omega^{\pi_i})^{-r(i)} g_i(\omega^{\pi_i}) \right).
\end{aligned}$$

Because the remaining sum is a determinant,

$$\left| \sum_{\pi \in \mathcal{S}_n} \det(P_\pi) \left(\prod_{i=1}^{n} (\omega^{\pi_i})^{-r(i)} g_i(\omega^{\pi_i}) \right) \right| =$$

$$\left| \det \begin{pmatrix} (\omega^0)^{-r(1)} g_1(\omega^0) & \cdots & (\omega^0)^{-r(n)} g_n(\omega^0) \\ (\omega^1)^{-r(1)} g_1(\omega^1) & \cdots & (\omega^1)^{-r(n)} g_n(\omega^1) \\ (\omega^2)^{-r(1)} g_1(\omega^2) & \cdots & (\omega^2)^{-r(n)} g_n(\omega^2) \\ \vdots & \cdots & \vdots \\ (\omega^{n-1})^{-r(1)} g_1(\omega^{n-1}) & \cdots & (\omega^{n-1})^{-r(n)} g_n(\omega^{n-1}) \end{pmatrix} \right|,$$

and because $\det(AB) = \det(A)\det(B)$,

$$\hat{f}(\chi) = n^{n/2} \left| \det \left(F \begin{pmatrix} (\omega^0)^{-r(1)} g_1(\omega^0) & \cdots & (\omega^0)^{-r(n)} g_n(\omega^0) \\ (\omega^1)^{-r(1)} g_1(\omega^1) & \cdots & (\omega^1)^{-r(n)} g_n(\omega^1) \\ (\omega^2)^{-r(1)} g_1(\omega^2) & \cdots & (\omega^2)^{-r(n)} g_n(\omega^2) \\ \vdots & \cdots & \vdots \\ (\omega^{n-1})^{-r(1)} g_1(\omega^{n-1}) & \cdots & (\omega^{n-1})^{-r(n)} g_n(\omega^{n-1}) \end{pmatrix} \right) \right|.$$

Now from (4) the main observation follows:

$$\hat{f}(\chi) = n^n \left| \det(L[r(i) + 1]) \right| \tag{5}$$

where L is the $n \times n$ matrix defined in section 2 and $L[\beta]$ is L with the approriate row rotations.

5 Bounds from the Uncertainty Principle

The uncertainty principle states roughly that, "a nonzero function and its transform cannot both be highly concentrated." Applied to Fourier analysis over a finite group, this principle has been interpreted in this way. Let G be a finite abelian group and f a nonzero function $f : G \to \mathbb{C}$. Donaho and Stark [4] proved that

$$\text{supp}(f)\text{supp}(\hat{f}) \geq |G|, \tag{6}$$

where $\text{supp}(f)$ denotes the support of the function f, and $\hat{f}$ is the transform of f. Matolcsi and Szűcs [8] (see also Diaconis and Shashahani [3]) obtained a similar formula for a compact group G and a nonzero mapping f, namely,

$$\text{supp}(f) \left(\sum_{\rho} \dim^2 \rho \right) \geq |G|,$$

where the sum is taken over irreducible representations ρ of G that have a nonzero coefficient in the unitary representation of f.

In this section we exploit the connection established between opening settings of the lock L_G and the Fourier transform of f_G to give bounds on the number of perfect matchings in G. Note that computing the number of perfect matchings in a bipartite graph is a $\#P$-complete problem as shown by Valiant ([12]) and is equivalent to a permanent computation. Thus providing bounds on the number of perfect matchings amounts to bounding a permanent.

There are other known bounds for the number of perfect matchings in certain classes of graphs, such as the bound proven by Voorhoeve in [13] for cubic bipartite graphs which was improved by Schrijver in [10] for k-regular bipartite graphs on $2n$ vertices. The bound in this paper is different because it works for

all bipartite graphs that have at least one perfect matching. However, finding a bound is not the main focus of this paper. In proving the main result of this paper it shows how the Combinatorial Nullstellensatz and the discrete Fourier transform can be used together to solve problems for which the Combinatorial Nullstellesatz seems applicable, such as bounding the number of perfect matchings in bipartite graphs. The hope of the authors is that similar methods will work for the problems discussed in the first section.

Once again, consider a bipartite graph $G(A, B; E)$ with vertex set $A \cup B$, where $A = \{0, \ldots, n-1\}$ and $B = \Omega_n$, and edge set

$$E \subseteq \{\{a, b\} : a \in A, b \in B\}$$

whose elements are the *edges* of G.

Theorem 4 *Let L be a circular lock derived from a bipartite graph G. If G has at least one perfect matching, then the product of the number of perfect matchings in G times the number of rotations that open the lock L is at least n^n.*

Proof. Apply the uncertainty principle for abelian groups to the function f_G and its transform over the group Ω_n^n. Simply observe that the definition (1) of f_G gives

$$\text{supp}(f_G) = \text{ number of perfect matchings of } G,$$

and observation (5) shows that

$$\text{supp}(\hat{f}_G) = \text{ number of rotations that open } L.$$

Because Ω_n^n is an abelian group of order n^n, the result follows from Donaho and Stark's version of the uncertainty principle (6). ◇

Let $m(G)$ denote the number of perfect matchings in the bipartite graph G and $r(G)$ the number of rotations that open the lock L_G. Theorem 4 states that $m(G)r(G) \geq n^n$, provided that $m(G) \neq 0$. Clearly equality is achieved when G is just a perfect matching. As another example, if G is the bipartite graph in Figure 2, then $m(G) = 2$ and $r(G) = 192$ and $2 \times 192 = 384 \geq 256 = 4^4$.

6 Determinants and Rotations

While working on the results presented in the previous sections, a related side result was discovered. In this section we investigate linear dependencies among the determinants

$$\{\det(A[\beta]) : \beta \in \mathbb{Z}_n^n\},$$

for a generic matrix A in which each cell contains a variable distinct from those in other cells. Our main theorem states that, for any vector $\alpha \in \mathbb{Z}_n^n$, the determinant of $A[\alpha]$ is an integer linear combination of the $n!$ determinants in the set

$$\{\det(A[\beta]) : \beta \in Z_n^n : \beta_i < n - i, \text{ for all } i = 0, \ldots n-1\}.$$

In particular, this shows that, to verify whether the polynomial given in (1) is in the ideal $\mathcal{I}_n$, one need not compute all of the coefficients, but rather at most $n!$ determinants.

Define $\mathcal{J}_n$ to be the ideal in $\mathbb{C}[x_0, \ldots, x_{n-1}]$ consisting of those polynomials that vanish on distinct nth roots of unity; that is, $f \in \mathcal{J}_n$ if and only if $f(\omega_0, \ldots, \omega_{n-1}) = 0$ for all $(\omega_0, \ldots, \omega_{n-1}) \in \Omega_n^n$ satisfying $\omega_i \neq \omega_j$, for $0 \leq i < j < n$.

Lemma 1 *For any $g = \sum_\alpha c_\alpha x^\alpha \in \mathcal{J}_n$ and any $n \times n$ matrix A,*

$$\sum_\alpha c_\alpha \det(A[\alpha]) = 0.$$

Proof. Suppose that

$$A = \begin{bmatrix} a_{0,0} & a_{0,1} & \cdots & a_{0,n-1} \\ a_{1,0} & a_{1,1} & \cdots & a_{1,n-1} \\ \vdots & \vdots & & \vdots \\ a_{n-1,0} & a_{n-1,1} & \cdots & a_{n-1,n-1} \end{bmatrix}$$

Now consider the ring homomorphism

$$h : \mathbb{C}[a_{0,0}, \ldots, a_{n-1,n-1}] \to \mathbb{C}[x_0, \ldots, x_{n-1}]$$

that maps $h(a_{i,j}) = x_i^j$. Let us write $h(A)$ for the matrix $[h(a_{i,j})]$. Clearly $h(A)$ is the Vandermonde matrix V shown below:

$$V = \begin{bmatrix} x_0^0 & x_0^1 & \cdots & x_0^{n-1} \\ x_1^0 & x_1^1 & \cdots & x_1^{n-1} \\ \vdots & \vdots & & \vdots \\ x_{n-1}^0 & x_{n-1}^1 & \cdots & x_{n-1}^{n-1} \end{bmatrix}$$

Recall that $\mathcal{I}_n$ is the ideal $\langle x_0^n - 1, \ldots, x_{n-1}^n - 1 \rangle$ in $\mathbb{C}[x_0, \ldots, x_{n-1}]$ of polynomials that vanish on all inputs from Ω_n^n. From elementary properties of determinants, we find that $\det(h(A[\alpha])) = x^\alpha \det(V) \pmod{\mathcal{I}_n}$. It follows that

$$\textstyle\sum_\alpha c_\alpha \det(h(A[\alpha])) = g \det(V) \pmod{\mathcal{I}_n}.$$

Now it is well known that $\det(V) = \prod_{0 \leq i \leq j < n} (x_j - x_i)$, so this determinant vanishes whenever there is some x_i with the same value as some x_j. Hence $g \det(V)$ vanishes on all inputs from Ω_n^n. Consequently $g \det(V)$ is the zero polynomial modulo $\mathcal{I}_n$. Because $\det(h(A[\alpha])) = h(\det(A[\alpha])) \pmod{\mathcal{I}_n}$, we deduce that

$$h\left(\textstyle\sum_\alpha c_\alpha \det(A[\alpha])\right) = 0 \pmod{\mathcal{I}_n}.$$

Now observe that each monomial of $\sum_\alpha c_\alpha \det(A[\alpha])$ contains at most one variable from each row of A, so $h\,(\sum_\alpha c_\alpha \det(A[\alpha])) = 0 \pmod{\mathcal{I}_n}$ implies actually that $\sum_\alpha c_\alpha \det(A[\alpha]) = 0$, as desired. ◇

Define, for $1 \le k \le n-1$,

$$f_k = \sum_{\alpha_0+\cdots+\alpha_k=n-k} x_0^{\alpha_0} \cdots x_k^{\alpha_k}$$

and $f_0 = x_0^n - 1$. The following fact is proven in [6]:
FACT: $\mathcal{J}_n = \langle f_0, f_1, \ldots, f_{n-1} \rangle$.

Let

$$R_n = \{\alpha \in \mathbb{Z}_n^n : \alpha_i < n-i, \text{ for all } i = 0, \ldots, n-1\}.$$

We now seek to define an ordering of the elements in $\mathbb{Z}_n^n$. First define, for any $\delta \in \mathbb{Z}_n^n \setminus R_n$,

$$I(\delta) = \max_{0 \le i \le n-1} \{i : \delta_i \ge n-i\}.$$

For two distinct $\delta, \gamma \in \mathbb{Z}_n^n$, define $\delta \le \gamma$ if

- $\delta \in R_n$, or
- $I(\delta) < I(\gamma)$, or
- $I(\delta) = I(\gamma)$ and $\delta_{I(\delta)} \le \gamma_{I(\gamma)}$.

As usual, we write $\delta < \gamma$, if $\delta \le \gamma$ and $\gamma \not\le \delta$. The significance of this ordering is the following lemma.

Lemma 2 *Suppose that A is an $n \times n$ matrix. If $\delta \in \mathbb{Z}_n^n \setminus R_n$, then $\det(A[\delta])$ can be expressed as a sum of the form $\sum_{\gamma<\delta} c_\gamma \det(A[\gamma])$, for some integer constants c_γ.*

Proof. Let $k = I(\delta)$. Define,

$$F_k = \{\alpha \in \mathbb{Z}_n^n : \sum_{i=0}^{k} \alpha_i = n-k \text{ and } \alpha_{k+1} = \cdots = \alpha_{n-1} = 0\}.$$

so F_k consists of the exponents of terms appearing in f_k. Define also

$$\epsilon = (0, \ldots, 0, n-k, 0, \ldots, 0),$$

where $n-k$ occurs in the kth coordinate. Clearly $\epsilon \in F_k$. Now set $\beta = \delta - \epsilon$. Since $f_k \in \mathcal{I}_n$, Lemma 1 shows that, for any matrix C,

$$\sum_{\alpha \in F_k} \det(C[\alpha]) = 0.$$

In particular, when $C = A[\beta]$, we deduce

$$\sum_{\alpha \in F_k} \det(A[\beta][\alpha]) = 0$$

where $A[\beta][\alpha]$ denotes the matrix obtained from A by first rotating β then further rotating α. Because $\det(A[\beta][\alpha]) = \det(A[\beta + \alpha])$ and $\delta = \beta + \epsilon$, we find

$$\det(A[\delta]) + \sum_{\alpha \in F_k \setminus \{\epsilon\}} \det(A[\beta + \alpha]) = 0.$$

It suffices now to observe that $\beta + \alpha < \delta$, for all $\alpha \in F_k \setminus \{\epsilon\}$. ◇

Now we are ready for the main theorem.

Theorem 5 *For any $n \times n$ matrix A and any $\beta \in \mathbb{Z}_n^n$, the determinant of $A[\beta]$ is an integer linear combination of the $n!$ determinants in the set*

$$\mathcal{P}_n = \{\det(A[\alpha]) : \alpha \in \mathbb{Z}_n^n : \alpha_i < n - i, \textit{ for all } i = 0, \ldots, n-1\}.$$

Proof. Use Lemma 2 to express $\det(A[\beta])$ as a "smaller" sum of the form $\sum_{\gamma < \beta} c_\gamma \det(A[\gamma])$, for some integer constants c_γ. Now repeatedly apply the lemma to replace any determinant in this sum with "smaller" sums until all determinants are in the set shown. ◇

It is not hard to see that Theorem 5 is best possible in the sense that the $n!$ multilinear polynomials in the set $\mathcal{P}_n$ are independent because the main diagonal of each matrix $A[\alpha]$ produces a monomial that does not appear in any other polynomial in this set.

7 Future Work

There are several directions for further research: 1). relate $r(G)$ to structural properties of G, 2). characterize the matroids that arise as the partition matroids and vector matroids of L_G as in the proof of Theorem 3, 3). extend these results to general matching and f-factor theorems, and 4). identify properties of these extensions that permit easy application of the Combinatorial Nullstellensatz.

References

[1] N. Alon, Combinatorial Nullstellensatz, *Combin. Probab. Comput.* 8 (1999), no. 1-2, 7–29.

[2] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, A. Schrijver, *Combinatorial optimization*, John Wiley & Sons, Inc., New York, 1998.

[3] P. Diaconis, M. Shashahani, On square roots of the uniform distribution on compact groups. *Proc. Amer. Math. Soc.* 98 (1986), no. 2, 341–348.

[4] D. L. Donoho, P. B. Stark, Uncertainty principles and signal recovery, *SIAM J. Appl. Math.* 49 (1989), no. 3, 906–931.

[5] A. E. Kézdy, ρ-valuations for some stunted trees, *Discrete Math.* 306 (2006), no. 21, 2786–2789.

[6] A. E. Kézdy, H. S. Snevily, Polynomials that vanish on distinct nth roots of unity. *Combin. Probab. Comput.* 13 (2004), no. 1, 37–59.

[7] E. Lawler, *Combinatorial optimization. Networks and matroids.* Reprint of the 1976 original. Dover Publications, Inc., Mineola, NY, 2001.

[8] T. Matolcsi, J. Szűcs, Intersection des mesures spectrales conjugues. (French) *C. R. Acad. Sci. Paris Sr. A-B* 277 (1973), A841–A843.

[9] J. G. Oxley, *Matroid theory*, Oxford University Press, New York, 1992.

[10] A. Schrijver, Counting 1-Factors in Regular Bipartite Graphs, *J. Combin. Theory, Ser. B* 72 (1998) no. 1, 122–135.

[11] A. Terras, *Fourier analysis on finite groups and applications* London Mathematical Society Student Texts, 43. Cambridge University Press, Cambridge, 1999.

[12] L. G. Valiant, The complexity of computing the permanent, *Theoret. Comput. Sci.* 8 (1979), no. 2, 189–201.

[13] M. Voorhoeve, A lower bound for the permanents of certain (0, 1)-matrices, *Nederl. Akad. Wetensch. Indag. Math.* 41 (1979) no. 1, 83–86.