Unlocking Matrices

Alexander Milner*

School of Mathematics, University of Edinburgh

Abstract

In this paper, we generalise a result from Section 2 of Brauch, Kézdy and Snevily's paper, [BKS14], where they first present the connection between bipartite graphs and when we can rotate the rows of a matrix so that it becomes invertible. They present this idea as an algorithm for determining whether a bipartite graph has a perfect matching by turning the problem into a question about matrices, which, in turn, can be solved in polynomial time using Edmond's Matroid Intersection Algorithm. However, by defining the bipartite graph first, they only consider a small selection of matrices with entries in $\mathbb{C}$. In the following, we extend these ideas to work for any matrix and over any field and we give an exact condition on when a matrix can be made invertible by rotating its rows. We then introduce the notions of cluster, minimal clusters and cluster density derived from the notion of the deficiency of a bipartite graph and use these to give an exact condition on when n^2 elements of a field can form an invertible $n \times n$ matrix.

1 Set-up

1.1 Notation

For $n \in \mathbb{N}$ and $\mathbb{F}$ a field, we let $[n] \coloneqq \{1, 2, ..., n\}$, $\mathbb{Z}_n \coloneqq \mathbb{Z}/n\mathbb{Z}$, S_n denote the set of permutations of $[n]$, $\overline{\mathbb{F}}$ denote the algebraic closure of $\mathbb{F}$ and μ_n denotes the set of the nth roots of unity, generated by a primitive nth root of unity ω.

When working with polynomials, $x \coloneqq (x_1, ..., x_n)$ and for $a \in \mathbb{F}^n$, then $x + a \coloneqq (x_1 + a_1, ..., x_n + a_n)$. Additionally, for $\alpha \in \mathbb{N}_0^n$, then $|\alpha| = \sum_{i=1}^n \alpha_i$ and $x^\alpha \coloneqq \prod_{i=1}^n x_i^{\alpha_i}$. Finally, for a polynomial $g \in \mathbb{F}[x]$, $Z(g)[\overline{\mathbb{F}}] \coloneqq \{x \in \overline{\mathbb{F}} : g(x) = 0\}$ is the set of roots of g from the algebraic closure of $\mathbb{F}$.

We denote a graph $G = (V(G), E(G)) = (V, E)$ where $V(G)$ is the set of vertices and $E(G)$ is the set of edges and for a subset of the vertices $W \subseteq V(G)$, $N_G(W) \subseteq V(G)$ denotes the set of neighbours of elements in W. If G is a bipartite graph, the vertices of G can be divided into two disjoint sets A and B, denoted $V(G) = (A, B)$, and we denote edges of G by (a, b) where $a \in A$ and $b \in B$.

1.2 Vandermonde's Matrix

Vandermonde's identity, a formula for the determinant of the Vandermonde matrix, crops up as a useful tool in a number of different areas of Combinatorics. As covered in Section 9.2 of [TV06], the determinant of the Vandermonde matrix is used in proofs of Dyson's conjecture and Snevily's Conjecture.

Definition 1.1 (Vandermonde matrix)**.** *For each $n \in \mathbb{N}$, let $V_n \in M_n(\mathbb{F}[x_1, ..., x_n])$ be the Vandermonde matrix, with elements $(V_n)_{i\,j} = x_i^{j-1}$ for all $i, j \in [n]$.*

*Email: A.J.C.Milner@sms.ed.ac.uk

Lemma 1.2 (Vandermonde's identity)**.** *If $V_n \in M_n(\mathbb{F}[x_1, ..., x_n])$ is the Vandermonde matrix, then*

$$\det V_n(x) = \prod_{1 \leq i < j \leq n} (x_j - x_i)$$

1.3 Combinatorial Nullstellensatz

The following is an incredibly useful result, commonly known as the Combinatorial Nullstellensatz, given as Theorem 1.1 in [Alo99].

Theorem 1.3. *Given an arbitrary field $\mathbb{F}$, let $f \in \mathbb{F}[x_1, ..., x_n]$. Let $S_1, ..., S_n$ be non-empty subsets of $\mathbb{F}$ and, for $i \in [n]$, define $g_i(x_i) \coloneqq \prod_{a \in S_i}(x_i - a) \in \mathbb{F}[x_i]$. If $f(s) = 0$ for all $s \in S_1 \times ... \times S_n$, then there exist polynomials $h_1, h_2, ..., h_n \in \mathbb{F}[x_1, ..., x_n]$, where, for all $i \in [n]$, $\deg(h_i) \leq \deg(f) - |S_i|$ such that $f = \sum_{i \in [n]} h_i g_i$.*

Later on we will need the following almost identical result to Theorem 1.3, which follows easily as a Corollary from it.

Corollary 1.4. *Given an arbitrary field $\mathbb{F}$, let $f \in \mathbb{F}[x_1, ..., x_n]$. Let $S_1, ..., S_n$ be non-empty subsets of $\mathbb{F}$ and, for $i \in [n]$, define $g_i(x_i) \coloneqq \prod_{a \in S_i}(x_i - a) \in \mathbb{F}[x_i]$. Then, $f(s) = 0$ for all $s \in S_1 \times ... \times S_n$ if and only if $f \in \langle g_i(x_i) : i \in [n] \rangle$.*

2 When can a matrix be unlocked...

Our original motivation is to look at the matrices given by rotating the rows of a given starting matrix and decide whether any of them are invertible. Since rotating the rows of a matrix is a group action on the matrix, it makes sense to generalise this immediately.

Definition 2.1. *Let the symmetric group S_{n^2} act on the n^2 elements of a matrix $M \in M_n(\mathbb{F})$ by permutation. Then, M is unlocked by a set $S \subseteq S_{n^2}$ if we can apply a sequence of group elements from S to M after which M is invertible ie. M is unlocked by S if $\exists\, \sigma \in \langle S \rangle \subseteq S_{n^2}$ such that $\det(\sigma(M)) \neq 0$.*

2.1 ...by rotations of its rows?

We can now restrict ourselves to cyclically permuting (or rotating) the rows of a matrix $M \in M_n(\mathbb{F})$.

Notation 2.2. *For $\alpha \in \mathbb{Z}_n^n$, let $M[\alpha]$ be the matrix defined by $(M[\alpha])_{i\,j} \coloneqq M_{(i \bmod n)\,(j+\alpha_i \bmod n)}$ ie. the matrix where we rotate the ith row to the left by α_i positions. Letting e_i denote the standard ith basis vector, if we let $r_i(M) \coloneqq M[e_i]$ then $r_i \in S_{n^2}$ and for $R \coloneqq \{r_i : i \in [n]\} \subseteq S_{n^2}$, then $\langle R \rangle = \{\prod_{i=1}^n r_i^{\alpha_i} : \alpha \in \mathbb{Z}_n^n\} \subseteq S_{n^2}$ since all r_i commute.*

In the language of Definition 2.1, we say M is unlocked by R or just M can be unlocked by rotations of its rows if and only if $\exists\, \sigma \in \langle R \rangle$ such that $\det(\sigma(M)) \neq 0$. This is equivalent to the existence of $\alpha \in \mathbb{Z}_n^n$ such that $\det(M[\alpha]) \neq 0$.

Example 2.3. *The matrix $\pi = \begin{pmatrix} 3 & -1 & -4 \\ 1 & 5 & -9 \\ 2 & -6 & 5 \end{pmatrix} \in M_3(\mathbb{Q})$ can be unlocked by row rotations as even though $\det(\pi) = 0$, $r_3(\pi) = \pi[e_3] = \begin{pmatrix} 3 & -1 & -4 \\ 1 & 5 & -9 \\ -6 & 5 & 2 \end{pmatrix}$ has determinant -27. $\gamma = \begin{pmatrix} 2 & -7 & 5 \\ -3 & -8 & 11 \\ 2 & 0 & -2 \end{pmatrix} \in M_3(\mathbb{Q})$,*

however, can not be unlocked by row rotations ie. $\det(\sigma(\gamma)) = 0$ *for all* $\sigma \in \langle R \rangle$*. This is simply due to the fact that the digits in each row of* γ *add up to* 0*. An easy way to see that this is the case is that* $(1,1,1)$ *is an eigenvector of* $\sigma(\gamma)$ *with eigenvalue* 0 *for all* $\sigma \in \langle R \rangle$*. Since the determinant of a matrix is equal to the product of its eigenvalues then* $\det(\sigma(\gamma)) = 0$ *for all* $\sigma \in \langle R \rangle$*. But, as we will see, the rows of a matrix all adding up to* 0 *is not a necessary condition for a matrix not to be able to be unlocked by row rotations.*

Definition 2.4. *[BKS14] Given a matrix* $M \in M_n(\mathbb{F})$*, for* $i \in [n]$*, define* $g_i \in \mathbb{F}[x_i]$ *by*

$$g_i(x_i) \coloneqq \sum_{j=1}^{n} M_{i\,j} x_i^{j-1}.$$

Letting V_n *be the* $n \times n$ *Vandermonde matrix as in Definition 1.1, define* $f_M \in \mathbb{F}[x_1, ..., x_n]$ *as*

$$f_M(x) \coloneqq (\det V_n)(x) \prod_{k=1}^{n} g_k(x_k) = \prod_{1 \le i < j \le n} (x_j - x_i) \prod_{k=1}^{n} g_k(x_k)$$

While the definition of our key polynomial f_M may look fairly arbitrary, its key features are that it contains all the information about our matrix M, both its elements and their positions, and also that it vanishes on any input $x = (x_1, ..., x_n)$ where $x_i = x_j$ for some $i \neq j$.

We now prove a technical lemma.

Lemma 2.5. *For permutations* $\sigma, \beta \in S_n$*, if* $\sigma(i) + \beta(i) \equiv k \pmod{n}$ *for all* $i \in [n]$*, then* $\operatorname{sgn}(\sigma) = (-1)^{(n-1)k + \lfloor \frac{n-1}{2} \rfloor} \operatorname{sgn}(\beta)$*.*

Proof. For permutations $\tau, \beta \in S_n$, we can restate the condition that $\tau(i) + \beta(i) \equiv 0 \pmod{n}$ for all $i \in [n]$ in terms of permutations by adding in transpositions as such

$$\tau = (1, n-1)(2, n-2)...(\lfloor \frac{n-1}{2} \rfloor, n - \lfloor \frac{n-1}{2} \rfloor)\beta = \left(\prod_{j=1}^{\lfloor \frac{n-1}{2} \rfloor} (j, n-j) \right) \beta$$

We need $\lfloor \frac{n-1}{2} \rfloor$ transpositions since for n odd, every element except n gets swapped ie. $\frac{n-1}{2} = \lfloor \frac{n-1}{2} \rfloor$ swaps whereas for n even, every element except n and $\frac{n}{2}$ gets swapped ie. $\frac{n-2}{2} = \lfloor \frac{n-1}{2} \rfloor$ swaps.

Finally, setting $\sigma = (1, ..., n)^k \tau$ implies $\sigma(i) \equiv \tau(i) + k \pmod{n}$ so $\sigma(i) + \beta(i) \equiv k \pmod{n}$ for all $i \in [n]$. Thus,

$$\begin{aligned} \operatorname{sgn}(\sigma) &= \operatorname{sgn}\left((1, ..., n)^k \prod_{j=1}^{\lfloor \frac{n-1}{2} \rfloor} (j, n-j)\beta \right) = \operatorname{sgn}(1, ..., n)^k \prod_{j=1}^{\lfloor \frac{n-1}{2} \rfloor} \operatorname{sgn}(j, n-j) \operatorname{sgn}(\beta) \\ &= (-1)^{(n-1)k + \lfloor \frac{n-1}{2} \rfloor} \operatorname{sgn}(\beta) \end{aligned}$$

since $\operatorname{sgn}((1, ..., n)) = (-1)^{n-1}$. □

The following Lemma is part of Theorem 2 from [BKS14], however, the factor of $(-1)^{\lfloor \frac{n-1}{2} \rfloor}$ is missed in the original paper which we amend here.

Lemma 2.6. *[BKS14] Given a matrix* $M \in M_n(\mathbb{F})$*, then*

$$f_M(x) \equiv (-1)^{\lfloor \frac{n-1}{2} \rfloor} \sum_{\alpha \in \mathbb{Z}_n^n} \det(M[\alpha]) x^{\alpha} \mod \langle x_i^n - 1 : i \in [n] \rangle$$

Proof. Working modulo the ideal $\langle x_i^n - 1 : i \in [n]\rangle$, we notice that

$$x^{-\alpha}\prod_{i=1}^{n} g_i(x_i) = \prod_{i=1}^{n} x_i^{-\alpha_i} g_i(x_i) = \prod_{i=1}^{n}\sum_{j=1}^{n} M_{i\,j} x_i^{j-1-\alpha_i} = \prod_{i=1}^{n}\sum_{j=1-\alpha_i}^{n-\alpha_i} (M)_{i\,j+\alpha_i} x_i^{j-1}$$
$$= \prod_{i=1}^{n}\sum_{j=1-\alpha_i}^{n-\alpha_i} (M[\alpha])_{i\,j\,(\mathrm{mod}\,n)} x_i^{j-1} \equiv \prod_{i=1}^{n}\sum_{j=1}^{n} (M[\alpha])_{i\,j} x_i^{j-1} = \sum_{\beta\in\mathbb{Z}_n^n}\prod_{i=1}^{n} x_i^{\beta_i-1}(M[\alpha])_{i\,\beta_i}$$

and using the Leibniz determinant formula

$$(\det V_n)(x) = \sum_{\sigma\in S_n}\mathrm{sgn}(\sigma)\prod_{k=1}^{n}(V_n)_{k\,\sigma(k)} = \sum_{\sigma\in S_n}\mathrm{sgn}(\sigma)\prod_{k=1}^{n} x_k^{\sigma(k)-1}$$

so

$$x^{-\alpha} f_M(x) \equiv \left(\sum_{\sigma\in S_n}\mathrm{sgn}(\sigma)\prod_{k=1}^{n} x_k^{\sigma(k)-1}\right)\left(\sum_{\beta\in\mathbb{Z}_n^n}\prod_{i=1}^{n} x_i^{\beta_i-1}(M[\alpha])_{i\,\beta_i}\right) \tag{1}$$

By comparing coefficients, we can see it is enough to show that, for any $\alpha \in \mathbb{Z}_n^n$, the constant term of $x^{-\alpha} f_M(x)$ modulo the ideal $\langle x_i^n - 1 : i \in [n]\rangle$ is equivalent to $(-1)^{\lfloor\frac{n-1}{2}\rfloor}\det(M[\alpha])$. Looking at Equation 1, the constant term of $x^{-\alpha} f_M(x)$ is given by the sum of $\mathrm{sgn}(\sigma)\prod_{i=1}^{n}(M[\alpha])_{i\,\beta_i}$ for $\sigma \in S_n, \beta \in \mathbb{Z}_n^n$ where $\sigma(k) - 1 \equiv -\beta_k - 1 \,(\mathrm{mod}\, n)$, equivalently $\sigma(k) \equiv -\beta_k \,(\mathrm{mod}\, n)$ for all $k \in [n]$. The only $\beta \in \mathbb{Z}_n^n$ that fulfil this are permutations of S_n since $\beta_k \equiv -\sigma(k) \,(\mathrm{mod}\, n)$ are distinct for all $k \in [n]$ since $\sigma \in S_n$. Using Lemma 2.5 with $k = 0$, the constant term of $x^{-\alpha} f_M(x)$ is thus

$$\sum_{\substack{\sigma\in S_n\\ \beta\in\mathbb{Z}_n^n\\ \sigma(k)\equiv-\beta_k}} sgn(\sigma)\prod_{i=1}^{n}(M[\alpha])_{i\,\beta_i} = \sum_{\substack{\sigma,\beta\in S_n\\ \sigma(i)+\beta(i)\equiv 0}} sgn(\sigma)\prod_{i=1}^{n}(M[\alpha])_{i\,\beta(i)}$$
$$= (-1)^{\lfloor\frac{n-1}{2}\rfloor}\sum_{\beta\in S_n}\mathrm{sgn}(\beta)\prod_{i=1}^{n}(M[\alpha])_{i\,\beta(i)} = (-1)^{\lfloor\frac{n-1}{2}\rfloor}\det(M[\alpha]).$$

□

Corollary 2.7. *Given a matrix $M \in M_n(\mathbb{F})$, then M is not unlocked by row rotations if and only if $f_M(x) \in \langle x_i^n - 1 : i \in [n]\rangle$.*

Remark 2.8. *Given a matrix M which can be unlocked by row rotations, Theorem 2.6 actually implies that calculating the polynomial expansion of $f_M(x) \bmod \langle x_i^n - 1 : i \in [n]\rangle$ automatically tells us which group elements in $\langle R\rangle \subseteq S_{n^2}$ it is possible to unlock the matrix for. Simply find those $\alpha \in \mathbb{Z}_n^n$ with non-zero coefficients in $f_M(x) \bmod \langle x_i^n - 1 : i \in [n]\rangle$ and then $\prod_{i=1}^{n} r_i^{\alpha_i}$ unlocks the matrix.*

We now have to consider two cases which depend on whether the character of our field $\mathbb{F}$ divides the size of our matrix n. These cases must be considered separately due to the differing number of nth roots of unity in $\mathbb{F}$ in each case.

Case 1: We now study matrices $M \in M_n(\mathbb{F})$ where $\mathrm{char}(\mathbb{F})|n$.

When $\mathrm{char}(\mathbb{F})|n$, we have $\mu_n = \{1\}$ as $(x-1)^n = \sum_{i=0}^{n}\binom{n}{i}(-1)^i x^{n-i} = x^n - 1$ since $n|\binom{n}{i}$ for $1 \le i \le n-1$ and $\mathrm{char}(\mathbb{F})|n$. Following Bruen in [Bru92], we define the multiplicity of an element in a polynomial over any field.

Definition 2.9. *For non-zero polynomial $g \in \mathbb{F}[x_1, ..., x_n]$, then $g(x)$ has multiplicity t at $a \in \mathbb{F}^n$ if*

$$t = \min\{|\alpha| : \alpha \in \mathbb{N}_0^n, c_\alpha \neq 0\} \quad \textit{where} \quad g(x+a) = \sum_{\alpha\in\mathbb{N}_0^n} c_\alpha x^\alpha$$

For convenience, if $g(x) = 0$, we say $g(x)$ has multiplicity ∞ for all $a \in \mathbb{F}^n$.

With Bruen's notion of multiplicity for multi-variable polynomials in hand, using Corollary 2.7, we can now prove an exact condition on when a matrix is unlocked by row rotations in the case that $\text{char}(\mathbb{F})|n$.

Theorem 2.10. *Let $M \in M_n(\mathbb{F})$ be a matrix where $\text{char}(\mathbb{F})|n$ with corresponding polynomials $g_i(x_i)$ for $i \in [n]$. Define the sequence $(n_i)_{i\in[n]}$, where $g_i(x_i)$ has multiplicity n_i at 1. Then, M is unlocked by row rotations if and only if there is a permutation $\sigma \in S_n$ such that $n_i < \sigma(i)$ for all $i \in [n]$.*

Proof. We start by proving $\forall\, \sigma \in S_n$, $\exists\, i \in [n]$ such that $n_i \geq \sigma(i) \Leftrightarrow f_M(x + 1_n) \in \langle x_i^n : i \in [n]\rangle$ where $1_n \coloneqq (1, ..., 1)$. For the $\Rightarrow$ direction, using the properties of the determinant of the Vandermonde matrix V_n from Lemma 1.2, we have

$$(\det V_n)(x + 1_n) = \prod_{1\leq i<j\leq n} (x_j + 1 - x_i - 1) = \prod_{1\leq i<j\leq n} (x_j - x_i) = (\det V_n)(x)$$

Also, from the definition of the n_i as multiplicities of the g_i, we have $g_i(x_i + 1) = h_i(x_i)x_i^{n_i}$ for some $h_i \in \mathbb{F}[x_i]$ where the h_i have a non-zero constant term. Thus, using the Leibniz determinant formula, we have

$$\begin{aligned} f_M(x + 1_n) &= (\det V_n)(x + 1_n) \prod_{i=1}^{n} g_i(x_i + 1) = (\det V_n)(x) \prod_{i=1}^{n} h_i(x_i)x_i^{n_i} \\ &= \left(\sum_{\tau\in S_n} \text{sgn}(\tau) \prod_{i=1}^{n} x_i^{\tau(i)-1} \right) \prod_{i=1}^{n} h_i(x_i)x_i^{n_i} = \sum_{\tau\in S_n} \text{sgn}(\tau) \prod_{i=1}^{n} x_i^{n_i-1+\tau(i)} h_i(x_i) \\ &= (-1)^{n-1+\lfloor \frac{n-1}{2} \rfloor} \prod_{i=1}^{n} h_i(x_i) \left(\sum_{\sigma\in S_n} \text{sgn}(\sigma) \prod_{i=1}^{n} x_i^{n+n_i-\sigma(i)} \right) \end{aligned}$$

where $\sigma \in S_n$ is defined by $\sigma(i) = (n + 1) - \tau(i)$ for all $i \in [n]$ and where we use Lemma 2.5 with $k = 1$ for the $(-1)^{n-1+\lfloor \frac{n-1}{2} \rfloor}$. It is now clear that if $\forall\, \sigma \in S_n$, $\exists\, i \in [n]$ such that $n_i \geq \sigma(i)$ then $f_M(x + 1_n) \in \langle x_i^n : i \in [n]\rangle$.

To prove the $\Leftarrow$ direction, we define a process. For some $Q \subseteq S_n$, define

$$f^{(Q)}(x) = (-1)^{n-1+\lfloor \frac{n-1}{2} \rfloor} \prod_{i=1}^{n} h_i(x_i) \left(\sum_{\sigma\in S_n\backslash Q} \text{sgn}(\sigma) \prod_{i=1}^{n} x_i^{n+n_i-\sigma(i)} \right)$$

and assume $f^{(Q)}(x) \in \langle x_i^n : i \in [n]\rangle$. Now let $d = \min\{\sum_{i\in[n]} n + n_i - \sigma(i) : \sigma \in S_n\backslash Q\}$ and $Q' = \{\sigma \in S_n\backslash Q : \sum_{i\in[n]} n + n_i - \sigma(i) = d\}$. Then since the h_i all have a non-zero constant term, the sum of monomials of $f^{(Q)}(x)$ with degree d is given by a multiple of $\sum_{\sigma\in Q'} \text{sgn}(\sigma) \prod_{i=1}^{n} x_i^{n+n_i-\sigma(i)}$. Since Q' is non-empty this sum is non-zero and thus $f^{(Q)}(x) \in \langle x_i^n : i \in [n]\rangle$ implies $\forall\, \sigma \in Q'$, $\exists\, i \in [n]$ such that $n_i \geq \sigma(i)$. In addition, $f^{(Q)}(x) \in \langle x_i^n : i \in [n]\rangle \Rightarrow f^{(Q\cup Q')}(x) \in \langle x_i^n : i \in [n]\rangle$ where $|Q \cup Q'| > |Q|$ and thus we can set $Q = Q \cup Q'$ and repeat the process.

We kick off the first iteration of this process by setting $Q = \emptyset$. Then since the size of Q strictly increases with each iteration, S_n is finite and $f_M(x + 1_n) = f^{(\emptyset)}(x) \in \langle x_i^n : i \in [n]\rangle$, we prove $\forall\, \sigma \in S_n$, $\exists\, i \in [n]$ such that $n_i \geq \sigma(i)$.

Forming a chain of equalities, $\forall\, \sigma \in S_n$, $\exists\, i \in [n]$ such that $n_i \geq \sigma(i) \Leftrightarrow f_M(x + 1_n) \in \langle x_i^n : i \in [n]\rangle$ $\Leftrightarrow f_M(x) \in \langle (x_i - 1)^n : i \in [n]\rangle = \langle x_i^n - 1 : i \in [n]\rangle \Leftrightarrow M$ is not unlocked by row rotations by Corollary 2.7. □

Example 2.11. $\begin{pmatrix} -1 & 1 & 0 \\ -1 & 0 & 1 \\ 0 & -1 & 1 \end{pmatrix} \in M_3(\mathbb{F}_3)$ *can not be unlocked by row rotations, since each row adds up to* 0 *and thus any polynomial with coefficients given by elements in a row will have* 1 *as a root so* $n_i \geq 1$ *for all* $i \in [n]$*. Then, for any* $\sigma \in S_n$*, taking* $\sigma(i) = 1$ *then* $n_i \not\leq \sigma(i) = 1$ *so by Theorem 2.10, the matrix can not be unlocked. Thus, it is easy to see that, in general, the rows adding up to* 0 *is a sufficient condition for a matrix to not be unlocked by row rotations. However, it is not a necessary condition. Take* $\begin{pmatrix} 1 & 2 & 0 & -1 & 0 \\ 1 & -1 & 2 & 0 & -2 \\ 2 & 1 & -1 & -2 & -2 \\ 2 & -1 & 0 & 1 & -2 \\ 1 & 0 & -1 & -2 & 2 \end{pmatrix} \in M_5(\mathbb{F}_5)$*, which cannot be unlocked by row rotations since the multiplicities at* 1 *are* $(0, 3, 0, 3, 3)$ *but rows* 1 *and* 3 *don't add up to* 0*.*

Case 2: We now study matrices $M \in M_n(\mathbb{F})$ where $\text{char}(\mathbb{F}) \nmid n$.

Lemma 2.12. *For* $n \in \mathbb{N}$*, let* $\mathbb{F}$ *be a field where* $\text{char}(\mathbb{F}) \nmid n$*. Then,* $t^n - 1 = \prod_{\tau \in \mu_n}(t - \tau)$*.*

Proof. Consider the roots of $x^n - 1$ in $\mathbb{F}$. Seeking a contradiction, assume $|\mu_n| < n$. Then $\exists\, \alpha \in \mu_n$ such that $x^n - 1 = (x - \alpha)^2 g(x)$. Taking the formal derivative on both sides, $D((x - \alpha)^2 g(x)) = (x - \alpha)^2 D(g(x)) + 2(x - \alpha)g(x) = D(x^n - 1) = nx^{n-1} =$ and plugging in α, we get $0 = n\alpha^{n-1} = \frac{n}{\alpha} \neq 0$ which is a contradiction. □

Lemma 2.13. *For* $n \in \mathbb{N}$*, let* $\mathbb{F}$ *be a field where* $\text{char}(\mathbb{F}) \nmid n$*. Then, given* $f \in \mathbb{F}[x_1, ..., x_n]$*,* $f \in \langle x_i^n - 1 : i \in [n] \rangle$ *if and only if* $\mu_n^n \subseteq Z(f)[\overline{\mathbb{F}}]$ *ie.* $f(x) = 0$ *for all* $x \in \mu_n^n$*.*

Proof. This follows by Corollary 1.4 using $g_i(x_i) = x_i^n - 1$ which factorises as $g_i(x_i) = \prod_{\omega \in \mu_n}(x_i - \omega)$ by Lemma 2.12. □

Definition 2.14. *Given a matrix* $M \in M_n(\mathbb{F})$*, where* $\text{char}(\mathbb{F}) \nmid n$*, with corresponding polynomials* $g_i(x_i)$*, define the bipartite graph* G_M *with vertices* $V(G_M) = ([n], \mu_n)$ *and edges given by* $(i, \omega^j) \in E(G_M)$ *if and only if* $\omega^j \notin Z(g_i)[\overline{\mathbb{F}}]$ *for all* $i, j \in [n]$*, where* ω *is a generator of* μ_n*.*

The following lemma is a generalisation of part of Theorem 2 from [BKS14] to any matrix over any field where $\text{char}(\mathbb{F}) \nmid n$.

Lemma 2.15. *Given a matrix* $M \in M_n(\mathbb{F})$ *with corresponding polynomial* $f_M(x)$ *and graph* G_M*, there exists a perfect matching on* G_M *if and only if* $\mu_n^n \not\subseteq Z(f_M)[\overline{\mathbb{F}}]$ *ie.* $f_M(x) \neq 0$ *for some* $x \in \mu_n^n$*.*

Proof. Let ω be a generator of μ_n. For the $\Rightarrow$ direction, let the perfect matching be given by $(i, \omega^{\sigma(i)}) \in E(G_M)$ for some $\sigma \in S_n$. Then by the definition of G_M, Definition 2.14, $\omega^{\sigma(i)} \notin Z(g_i)[\overline{\mathbb{F}}]$ for all $i \in [n]$. Thus since σ is a permutation and using Lemma 1.2, then $\det V_n(\omega^{\sigma(1)}, ..., \omega^{\sigma(n)}) \neq 0$, thus $f_M(\omega^{\sigma(1)}, ..., \omega^{\sigma(n)}) \neq 0$. For the $\Leftarrow$ direction, assume $f_M(x) \neq 0$ for some $x \in \mu_n^n$, then since $(\det V_n)(x) \neq 0$, again using Lemma 1.2, then $x = (\omega^{\sigma(1)}, ..., \omega^{\sigma(n)})$ for some $\sigma \in S_n$ a permutation. Thus, for all $i \in [n]$, $\omega^{\sigma(i)} \notin Z(g_i)[\overline{\mathbb{F}}] \Leftrightarrow (i, \omega^{\sigma(i)}) \in E(G_M)$ and since σ is a permutation, this is a perfect matching on G_M. □

The equivalent of Theorem 2.10 for the case $\text{char}(\mathbb{F}) \nmid n$ now falls out.

Theorem 2.16. *Given a matrix* $M \in M_n(\mathbb{F})$ *where* $\text{char}(\mathbb{F}) \nmid n$ *with corresponding graph* G_M*, then* M *is unlocked by row rotations* $\Leftrightarrow$ *there exists a perfect matching on* G_M*.*

Proof. By Corollary 2.7, Lemma 2.13 and Lemma 2.15, M is unlocked by row rotations $\Leftrightarrow f_M \notin \langle x_i^n - 1 : i \in [n] \rangle \Leftrightarrow \mu_n^n \not\subseteq Z(f_M)[\overline{\mathbb{F}}] \Leftrightarrow$ there exists a perfect matching on G_M. $\square$

Example 2.17. *We can now see another reason why the matrix* $\gamma = \begin{pmatrix} 2 & -7 & 5 \\ -3 & -8 & 11 \\ 2 & 0 & -2 \end{pmatrix} \in M_3(\mathbb{R})$ *from Example 2.3, or indeed any matrix where the rows add up to* 0 *can not be unlocked by row rotations. We can construct a polynomial with coefficients given by a row of the matrix; if the coefficients add up to* 0*, then the polynomial has a root at* 1*. If this is the case for every row, then when we construct our bipartite graph* G_M*,* $(i, 1) \notin E(G)$ *for any* $i \in [n]$*. Thus, by the pigeonhole principle, it is impossible for* G_M *to have a perfect matching, otherwise there would have to be an edge connected to* $1 \in \mu_n$*.*

For a more complicated example, let $\mathbb{F} = \mathbb{F}_3$*, and take* $\mu = \begin{pmatrix} 0 & 1 & 0 & 1 \\ -1 & -1 & -1 & -1 \\ -1 & 1 & 0 & -1 \\ 1 & -1 & 1 & -1 \end{pmatrix} \in M_4(\mathbb{F}_3)$*. You would have to calculate at most* $4! = 24$ *determinants to find out if* μ *can be unlocked by rotations, however, luckily for us we have some theorems we can use.*

Instead of taking the full closure of $\mathbb{F}_3$*, it is enough to let* $\theta^2 + 1 = 0$ *and work in* $\mathbb{F}_9 = \mathbb{F}_3[\theta]$ *since* $\mathbb{F}_9$ *contains* Ω_4 *as* $\mathbb{F}_9^\times \cong \mathbb{Z}_8$*. Then notice that polynomials* $g_1(x_1) = x_1^3 + x_1$*,* $g_2(x_2) = -x_2^3 - x_2^2 - x_2 - 1$ *and* $g_4(x_4) = -x_4^3 + x_4^2 - x_4 + 1$ *all have* θ *and* $-\theta$ *as roots. Thus* G_μ*, as depicted in Figure 1, does not have a perfect matching since only* 3 *is connected to* θ *and* $-\theta$ *and thus by Theorem 2.16,* μ *can not be unlocked by row rotations.*

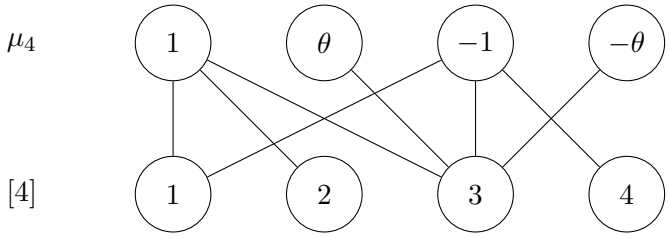

Figure 1: The bipartite graph G_μ corresponding to matrix μ from Example 2.17.

Remark 2.18. *It seems that whether* $\mathrm{char}(\mathbb{F})$ *divides* n *or not gives rather different conditions on when the matrix can be unlocked by row rotations. However, these two cases are not so dissimilar if we reformulate the case* $\mathrm{char}(\mathbb{F})|n$*. The constraint given in Theorem 2.10 was that* $M \in M_n(\mathbb{F})$ *was unlocked by row rotations if and only if there was a permutation* $\sigma \in S_n$ *such that* $n_i < \sigma(i)$ *for* $i \in [n]$ *where* $(n_i)_{i \in [n]}$ *was given by* $g_i(x_i)$ *having multiplicity* n_i *at* 1*. Similarly to Definition 2.14, where we defined* G_M*, we can define a graph* H_M *with vertices given by* $V(H_M) = ([n], [n])$ *and edges* $(i, j) \in E(H_M)$ *if and only if* $n_i < j$ *for all* $i, j \in [n]$*. Clearly* $\exists\, \sigma \in S_n$ *such that* $n_i < \sigma(i)$ *for all* $i \in [n]$ $\Leftrightarrow (i, \sigma(i)) \in E(H_M) \Leftrightarrow H_M$ *has a perfect matching since* σ *is a permutation. In some sense, we see that Theorem 2.10 (the* $\mathrm{char}(\mathbb{F})|n$ *case) is not as strong a statement as Theorem 2.16 (the* $\mathrm{char}(\mathbb{F}) \nmid n$ *case), since we can construct matrices* M *such that* G_M *is any bipartite graph whereas this is not true for* H_M*, since for any edge* $(i, j) \in E(H_M)$*, we automatically have* $(i, j') \in E(H_M)$ *for all* $j' \geq j$*. This will result in a more difficult proof of the more general versions of Theorem 2.10 and Theorem 2.16, where we want to prove the forms of matrices which are unlocked by the set of all permutations.*

2.2 ...by all permutations?

We now have an exact condition on when a matrix is unlocked by row rotations. But what happens if we allow ourselves to rotate both the rows and columns as we please in any order? What if we allow any permutation of the elements of the matrix? It turns out that the latter question will help us answer the former, so we tackle that first.

We will first give an exact condition on when a matrix is unlocked by all permutations for the easier $\text{char}(\mathbb{F})|n$ case, as a warm-up for the trickier $\text{char}(\mathbb{F}) \nmid n$ case.

Theorem 2.19. *For $n \geq 2$, given n^2 elements in a field $\mathbb{F}$ with $\text{char}(\mathbb{F})|n$, where there are at most $n^2 - n + 1$ of the same element or at most $n^2 - n$ zeroes, we can always construct an invertible $n \times n$ matrix out of those elements.*

Before we prove Theorem 2.19, we must first prove some technical lemmas.

Lemma 2.20. *For $n \geq 1$, given a polynomial $p \in \mathbb{F}[t]$ with degree at most $n-1$, $Z(p)[\overline{\mathbb{F}}] \cap \mu_n$ is invariant under cyclic permutations (rotations) of the coefficients of $p(t)$.*

Proof. The statement of this Lemma is equivalent to $\omega \in Z(p(t))[\overline{\mathbb{F}}] \Leftrightarrow \omega \in Z(t^k p(t) \bmod \langle t^n - 1 \rangle)[\overline{\mathbb{F}}]$ for all $k \in [n]$, $\omega \in \mu_n$. Writing, $t^k p(t) \bmod \langle t^n - 1 \rangle$as $t^k p(t) + (t^n - 1)q(t)$ for some $q \in \mathbb{F}[t]$, and evaluating at ω, we get $\omega^k p(\omega) + (\omega^n - 1)q(\omega) = \omega^k p(\omega)$ and the result follows since $\omega^k \neq 0, \forall\, k \in [n]$. □

Lemma 2.21. *For $n \geq 2$, given a polynomial $p \in \mathbb{F}[t]$, let $p(t) = \sum_{i=0}^{n-1} a_i t^i$. For some $b_0 \in \mathbb{F}$ where $b_0 \neq a_0$, let $\hat{p}(t) = \sum_{i=1}^{n-1} a_i t^i + b_0$. Then $p(t)$ and $\hat{p}(t)$ share no roots in $\overline{\mathbb{F}}$ ie. $Z(p)[\overline{\mathbb{F}}] \cap Z(\hat{p})[\overline{\mathbb{F}}] = \emptyset$.*

Proof. We have $p(t) - \hat{p}(t) = a_0 - b_0 \neq 0$, thus $p(t)$ and $\hat{p}(t)$ share no common values, in particular share no roots. □

Corollary 2.22. *Given a matrix M with corresponding sequence of multiplicities $(n_i)_{i=1}^n$, if for $i, j \in [n]$, $0 < n_i, n_j \leq n-1$, after swapping distinct elements of M, one from row i and one from row j, the new matrix will have multiplicities $n_i = n_j = 0$.*

Proof. Let α_i, α_j be the elements to be swapped in rows i, j respectively. Since $n_i, n_j > 0$, then $1 \in Z(g_i)[\overline{\mathbb{F}}], Z(g_j)[\overline{\mathbb{F}}]$. We start by rotating rows i, j, thus cyclically permuting the coefficients of g_i, g_j until α_i, α_j are the constant coefficients in g_i, g_j respectively ie. α_i, α_j are in the first column of M. By Lemma 2.20, 1 is still a root of g_i and g_j. Using Lemma 2.21 on both polynomials g_i, g_j separately, when we swap α_i, α_j, $1 \notin Z(g_i)[\overline{\mathbb{F}}], Z(g_j)[\overline{\mathbb{F}}]$. Finally, we can rotate rows i, j until α_i is in the old position of α_j and α_j is in the old position of α_i and, using 2.20, even after these rotations, $1 \notin Z(g_i)[\overline{\mathbb{F}}], Z(g_j)[\overline{\mathbb{F}}]$. The new matrix is just M with α_i and α_j swapped but since $1 \notin Z(g_i)[\overline{\mathbb{F}}], Z(g_j)[\overline{\mathbb{F}}]$ in the new matrix, $n_i = n_j = 0$. □

We will now prove Theorem 2.19. We start by arranging the n^2 elements in a matrix M (we can't guarantee M is invertible). We then swap a series of elements between rows until we can guarantee that the matrix is able to be unlocked by row rotations using our exact statement on when a matrix can be unlocked, Theorem 2.10. Finally, we can perform the relevant row rotations, leaving us with an invertible matrix made from the n^2 elements. Even though we are performing changes to the matrix M we will not keep track of these and will continually denote our matrix M.

Proof of Theorem 2.19. Since we have at most $n^2 - n + 1$ of the same element, we can always arrange the n^2 elements in a matrix M such that at most 1 row contains n copies of the same element and no

row contains all zeroes. The condition that at most 1 row contains n copies of the same non-zero element implies that there is at most $i \in [n]$ such that $n_i = n-1$. This is because the polynomial with all entries the same and non-zero is equal to a multiple of $x^{n-1} + x^{n-2} + ... + x + 1 = \frac{x^n-1}{x-1} = \frac{(x-1)^n}{x-1} = (x-1)^{n-1}$. The condition that no row contains all zeroes implies that every n_i is finite and since $\deg(g_i(x_i)) \leq n-1$, $n_i \leq n-1$ for all $i \in [n]$.

We will now swap distinct elements from distinct rows, until all but one $n_i = 0$ for $i \in [n]$. Given two rows $i, j \in [n]$ with multiplicities $0 < n_i, n_j \leq n-1$, we can always find two distinct elements, one from each row, since at most one row contains n copies of the same element. By Corollary 2.22, after swapping these elements, $n_i = n_j = 0$. It is important to note that, in doing this, we never create a row which contains n copies of the same element and thus we can always guarantee that at most one row of the matrix M has n copies of the same element.

We repeat this process until there is at most one row $i \in [n]$ with $n_i \neq 0$. It is clear no elements have been swapped in row i since otherwise $n_i = 0$. However, since $n_i \leq n-1$ to start off with, we can choose any $\sigma \in S_n$ with $\sigma(i) = n$ and since all other $n_j = 0$, $n_j < \sigma(j)$ for all $j \in [n]$. Thus, by Theorem 2.10, M is unlocked by row rotations. Thus, by applying the necessary swaps and row rotations to our starting matrix we constructed an invertible matrix out of our original n^2 elements. □

Unfortunately, due to the more complex condition involving perfect matchings on bipartite graphs for the case when $\text{char}(\mathbb{F}) \nmid n$, we will not be able to prove the same statement as in Theorem 2.19 straight away. Instead, we must state Hall's marriage, an exact condition on when bipartite graphs have perfect matchings, and introduce the key concepts of clusters, minimal clusters and cluster density.

Hall's Marriage Theorem

First proved by Hall in [Hal86], Hall's marriage theorem gives an exact condition on when a bipartite graph has a perfect matching. According to [Hir07], it supposedly got its name from one of the many ways the theorem can be posed: suppose we have a group of boys and girls, where we need to find all the boys a partner from the group of girls. We can start by asking the girls to write a list of the boys they find acceptable and we assume the boys will not turn down a date with a girl. Given this information, can we match the boys and girls up in happy couples?

Definition 2.23. *We recall that a perfect matching on a graph G is a subset of the edge set $S \subseteq E(G)$ such that every vertex in $V(G)$ is contained in some edge in S. Now let G be a bipartite graph with vertices $V(G) = (A, B)$. Then, we define an $A-$perfect matching on G to be a subset of the edge set $S \subseteq E(G)$, such that every vertex of A is contained in some edge in S.*

Remark 2.24. *For a bipartite graph G with vertices $V(G) = (A, B)$, if $|A| = |B|$, G has an $A-$perfect matching $\Leftrightarrow$ G has a $B-$perfect matching $\Leftrightarrow$ G has a perfect matching.*

We now state but do not prove Hall's Theorem. A good proof is found in [DeV] using the theory of M-alternating and M-augmenting paths.

Theorem 2.25 (Hall's marriage theorem)**.** *Given a bipartite graph G with vertices $V(G) = (A, B)$, there exists an $A-$perfect matching on $G \Leftrightarrow |W| \leq |N_G(W)|$ for all $W \subseteq A$.*

We now introduce some notions that tie in closely with Hall's Marriage Theorem and perfect matchings, the first being the deficiency of a bipartite graph, originally defined by Ore in [Ore55].

Definition 2.26. *Given a bipartite graph G with vertices $V(G) = (A, B)$, the deficiency of a set $U \subseteq V(G)$, is defined to be* $\mathrm{def}_G(U) \coloneqq |U| - |N_G(U)|$. *Furthermore, the deficiency of G with respect to A is defined to be* $\mathrm{def}(G; A) \coloneqq \max_{U \subseteq A} \mathrm{def}_G(U)$. *Note that* $\mathrm{def}_G(\emptyset) = 0$ *so we have that* $\mathrm{def}(G; A) \geq 0$. *Finally, if $|A| = |B|$, the deficiency of G is defined to be* $\mathrm{def}(G) \coloneqq \mathrm{def}(G; A) = \mathrm{def}(G; B)$.

Lemma 2.27. *For a bipartite graph G where $|A| = |B|$, then* $\mathrm{def}(G; A) = \mathrm{def}(G; B)$.

Proof. For a set $U \subseteq A$, by Definition 2.26, we have

$$\begin{aligned}\mathrm{def}_G(U) &= |U| - |N_G(U)| = |A| - |A \backslash U| - |B| + |B \backslash N_G(U)| \\ &\leq -|N_G(B \backslash N_G(U))| + |B \backslash N_G(U)| = \mathrm{def}_G(B \backslash N_G(U)).\end{aligned}$$

Thus $\mathrm{def}(G; A) \leq \mathrm{def}(G; B)$, so by symmetry of swapping A and B, $\mathrm{def}(G; A) = \mathrm{def}(G; B)$. □

We now define the original notions of clusters, minimal clusters and cluster density.

Definition 2.28. *Let G be a bipartite graph with vertices $V(G) = (A, B)$. We define a cluster in G to be a set $W \subseteq A$ where* $\mathrm{def}_G(W) > 0$ *ie.* $|W| > |N_G(W)|$. *Furthermore we say that a cluster $W \subseteq A$ is minimal if there does not exist a set $U \subset W \subseteq A$ which is again a cluster. Let* $\mathrm{clust}(G; A) = \{W \subseteq A : \mathrm{def}_G(W) > 0\}$ *ie. the set of clusters in G and define the cluster density to be* $\mathrm{cd}(G; A) \coloneqq \sum_{W \in \mathrm{clust}(G;A)} \mathrm{def}_G(W)$.

Remark 2.29. *It is easy to see, using Hall's Marriage Theorem, Theorem 2.25, and Definition 2.26 and Definition 2.28, for a bipartite graph G with vertices $V(G) = (A, B)$, the following are equivalent:*

- *G has an $A-$perfect matching,*
- *$|W| \leq |N_G(W)|$ for all $W \subseteq A$,*
- *$\mathrm{def}_G(W) \leq 0$ for all $W \subseteq A$,*
- *$\mathrm{def}(G; A) = 0$,*
- *$\mathrm{cd}(G; A) = 0$.*

We can finally return to our study of matrices and state the equivalent of Theorem 2.19 but now for the case $\mathrm{char}(\mathbb{F}) \nmid n$, which has a similar but substantially harder proof. We will then be able combine it with Theorem 2.19 to give an exact statement on when all matrices are unlocked by all permutations.

Theorem 2.30. *For $n \geq 3$, given n^2 elements in a field $\mathbb{F}$ with* $\mathrm{char}(\mathbb{F}) \nmid n$, *where there are at most $n^2 - n + 1$ of the same element or at most $n^2 - n$ zeroes, we can always construct an invertible $n \times n$ matrix out of those elements.*

Before we prove Theorem 2.30, we will need some technical lemmas. These aim to show that, by swapping elements of a matrix M, we can strictly reduce the cluster density of the corresponding bipartite graph G_M and, in some sense, remove clusters until we are guaranteed to be left with a bipartite graph which has a perfect matching, meaning M can be unlocked by row rotations.

Definition 2.31. *Let G be a bipartite graph with vertices $V(G) = (A, B)$. We will now define a set of graphs, depending on some point $p \in A$ and denoted $T_p(G)$ which we refer to as transformed graphs. A graph $G' \in T_p(G)$ if:*

- *$V(G') = (A, B)$ ie. G' has the same vertices as G,*
- *$(p, b) \notin G \Rightarrow (p, b) \in G'$ for all $b \in B$,*
- *$(q, b) \in G \Leftrightarrow (q, b) \in G'$ for all $q \in A \backslash \{p\}$.*

For a set of bipartite graphs S, let $T_p(S) \coloneqq \cup_{G \in S} T_p(G)$.

Lemma 2.32. *Let G be a bipartite graph with vertices $V(G) = (A, B)$ where $|A| = |B|$. Let W be a minimal cluster in G, then for any $p \in W$, for all $G' \in T_p(G)$, $\mathrm{cd}(G'; A) < \mathrm{cd}(G; A)$.*

Proof. Let $p \in W$ and $G' \in T_p(G)$. Then we aim to show that any set $W' \subseteq A$ where $p \in W'$ is not a cluster in G'.

- Letting $Q = (W \cap W')\backslash\{p\}$, then Q is not a cluster in G since otherwise W would not be minimal as $Q \subseteq W$. Thus $|N_G(Q)| \geq |Q|$.
- We now claim $N_G(Q) \sqcup (\mu_n \backslash N_G(W)) \subseteq N_{G'}(W')$. Since $(q, b) \in G \Leftrightarrow (q, b) \in G'$ for all $q \in A\backslash\{p\}$ and $p \notin Q$, then $N_G(Q) = N_{G'}(Q) \subseteq N_{G'}(W')$ as $Q \subseteq W'$. Furthermore, since $(p, b) \notin G \Rightarrow (p, b) \in G'$ for all $b \in B$ and $p \in W, W'$, we also have that $\mu_n = N_G(\{p\}) \cup N_{G'}(\{p\}) \subseteq N_G(W) \cup N_{G'}(W')$, and thus $(\mu_n \backslash N_G(W)) \subseteq N_{G'}(W')$ showing $N_G(Q) \cup (\mu_n \backslash N_G(W)) \subseteq N_{G'}(W')$. Finally, $N_G(Q) \cap (\mu_n \backslash N_G(W)) \subseteq N_G(W) \cap (\mu_n \backslash N_G(W)) = \emptyset$ and thus $|N_{G'}(W')| \geq |N_G(Q)| + |\mu_n \backslash N_G(W)|$.
- Since W is a cluster, we know $|W| > |N_G(W)|$, so as $W \subseteq [n]$ and $N_G(W) \subseteq \mu_n$, $|\mu_n \backslash N_G(W)| = n - |N_G(W)| > n - |W| = |[n]\backslash W| \geq |W'\backslash W|$. Thus $|\mu_n \backslash N_G(W)| \geq |W'\backslash W| + 1$.

Since $|N_{G'}(W')| \geq |N_G(Q)| + |\mu_n \backslash N_G(W)| \geq |Q| + |W'\backslash W| + 1 = |W \cap W'| - 1 + |W'\backslash W| + 1 = |W'\backslash W| + |W' \cap W| = |W'|$ then W' is not a cluster. Finally, let $U \in \mathrm{clust}(G'; A)$. By the above, $p \notin U$, so since $(q, b) \in G \Leftrightarrow (q, b) \in G'$ for all $q \in A\backslash\{p\}$ then $\mathrm{def}_{G'}(U) = |U| - |N_{G'}(U)| = |U| - |N_G(U)| = \mathrm{def}_G(U)$. Also, since $p \in W$, $W \notin \mathrm{clust}(G'; A)$ but $W \in \mathrm{clust}(G; A)$ and thus $\mathrm{cd}(G'; A) = \sum_{U \in \mathrm{clust}(G'; A)} \mathrm{def}_{G'}(U) < \sum_{U \in \mathrm{clust}(G; A)} \mathrm{def}_G(U) = \mathrm{cd}(G'; A)$. □

Remark 2.33. *It is interesting to note that Lemma 2.32 would not necessarily hold if we just required the point p to be in a cluster and not a minimal cluster.*

We now need to build up some results about how swapping elements in our matrix M affects the corresponding bipartite graph G_M

Lemma 2.34. *For $n \geq 2$, given a polynomial $p \in \overline{\mathbb{F}}[t]$, let $p(t) = \sum_{i=0}^{n-1} a_i t^i$. Consider the polynomial where we swap the first two entries ie. $\hat{p}(t) = \sum_{i=2}^{n-1} a_i t^i + a_0 t + a_1$. Then, if $a_0 \neq a_1$, $p(t)$ and $\hat{p}(t)$ share no roots except for $t = 1$ ie. $Z(p)[\overline{\mathbb{F}}] \cap Z(\hat{p})[\overline{\mathbb{F}}] \in \{\emptyset, \{1\}\}$.*

Proof. The result follows by considering $p(t) - \hat{p}(t) = (a_1 - a_0)(t - 1)$ and using $a_0 \neq a_1$. Thus, $p(t)$ and $\hat{p}(t)$ share no common values, in particular no common roots, unless $t = 1$. □

Lemma 2.35. *Given a matrix M and bipartite graph G_M, if M' is the matrix where we replace any element in row $i \in [n]$ with a different element, then $G_{M'} \in T_i(G_M)$.*

Proof. Let M have corresponding polynomials $g_i(x_i) = \sum_{i \in [n]} M_{ij} x_i^{j-1}$ for all $i \in [n]$ and let α be the element we want to replace which is in position (i, j). Let $\hat{M} \coloneqq r_i^{j-1}(M)$ ie. the matrix where we can rotate the elements in the ith row so that α is now in position $(i, 1)$ and thus acts as the constant of $g_i(x_i)$. By Lemma 2.20, this does not change $Z(g_i)[\overline{\mathbb{F}}] \cap \mu_n$, thus $G_{\hat{M}} = G_M$. Now, let $\overline{M}$ be the matrix where we replace α with α'. By Lemma 2.21, if $\overline{g}_i(x_i) = \sum_{i \in [n]} \overline{M}_{ij} x_i^j$, since $\alpha \neq \alpha'$ then $Z(g_i)[\overline{\mathbb{F}}] \cap Z(\overline{g}_i)[\overline{\mathbb{F}}] \cap \mu_n = \emptyset$, thus in $G_{\overline{M}}$ the vertex $i \in [n]$ is connected to all the vertices in μ_n that it wasn't connected to in G_M and all other vertices in $[n]$ and their edges are identical, thus $G_{\overline{M}} \in T_i(G_M)$. Finally, let M' be the matrix where we undo the rotation we did at the beginning so M' is simply M with one element in row i changed. By Lemma 2.20 again, $G_{\overline{M}} = G_{M'}$ which implies $G_{M'} \in T_i(G_M)$. □

Notation 2.36. *For a graph G, for $S \subseteq V(G)$, the induced subgraph $G[S]$ is the graph whose vertex set is S and whose edge set are those edges in G where both endpoints are in S ie. $V(G[S]) = S$ and*

$(i,j) \in E(G[S]) \Leftrightarrow (i,j) \in E(G)$ *and* $i,j \in S$. *Furthermore, if* G *is bipartite with vertices* $V(G) = (A,B)$, *then for* $S_A \subseteq A$ *and* $S_B \subseteq B$, *the induced subgraph is denoted* $G[(S_A, S_B)]$.

Corollary 2.37. *For* $A \subseteq [n]$, *if* $\tilde{M}$ *is the matrix where we swap two different adjacent elements in row* i *of* M, *then* $G_{\tilde{M}}[(A, \mu_n \backslash \{1\})] \in T_i(G_M[(A, \mu_n \backslash \{1\})])$.

Proof. By Lemma 2.20, we can move the two adjacent elements to be the first two row entries so that they are the constant and linear term in $g_i(x_i)$, leaving G_M unchanged. By Lemma 2.34, when we swap them, all the previous roots of $g_i(x_i)$ are no longer roots except for $x_i = 1$. Thus, the corresponding graph of the new matrix lies in $T_i(G_M[(A, \mu_n \backslash \{1\})])$. Finally, we use Lemma 2.20 to move the elements back to their starting positions implying $G_{\tilde{M}}[(A, \mu_n \backslash \{1\})] \in T_i(G_M[(A, \mu_n \backslash \{1\})])$. □

We are now in a position to prove Theorem 2.30, which will follow the same process we used for the $\text{char}(\mathbb{F})|n$, Theorem 2.19. We arrange our n^2 elements in a matrix M and perform a series of swaps of elements, thus reducing the cluster density of the corresponding bipartite graph G_M, until we can guarantee the matrix can be unlocked by row rotation using our exact condition on when matrices can be unlocked, Theorem 2.16. Applying the relevant row rotations, the matrix made out of the n^2 elements is now invertible.

Proof of Theorem 2.30. To start, since we have at most $n^2 - n + 1$ of the same element, we can always arrange the n^2 elements in the matrix M such that at most 1 row contains n copies of the same element and no row contains all zeroes. The condition that at most 1 row contains n copies of the same non-zero element implies that there is at most 1 vertex in $[n]$ which has only 1 edge and that edge is connected to $1 \in \mu_n$. This is because the polynomial with all entries the same and non-zero is equal to a multiple of $x^{n-1} + x^{n-2} + ... + x + 1 = \frac{x^n - 1}{x - 1}$ and thus has roots $\mu_n \backslash \{1\}$. The condition that no row contains all zeroes implies that every vertex in $[n]$ has at least 1 edge connected to it, since everything, including μ_n, is a root of the zero polynomial.

We need to resolve a few technicalities before we perform the majority of the switches. In particular, we need $1 \in \mu_n$ to be connected to at least 1 vertex in $[n]$ but still keep the condition that there is at most 1 vertex in $[n]$ which has only 1 edge and that edge is connected to $1 \in \mu_n$. If $1 \in \mu_n$ has no edges, we will swap two elements of M to rectify this. Since at most 1 row contains all the same elements, we can always find two different elements α and β in two different rows i and j respectively to swap, leaving us with a new matrix $\hat{M}$. By Lemma 2.35, since $(i,1), (j,1)$ are not edges of G_M, then $(i,1), (j,1)$ are edges in all the graphs in $T_i(T_j(G_M))$ and $G_{\hat{M}} \in T_i(T_j(G_M))$. There is now the unwanted case that in $G_{\hat{M}}$, $i, j \in [n]$ are both now only connected to $1 \in \mu_n$. If this happens, we know that rows i and j of $\hat{M}$ are both filled with n copies of β and α respectively. So swap α and β back so our matrix returns to M and now swap an extra β and α between rows i and j giving us a new matrix $\overline{M}$ with $G_{\overline{M}} \in T_i(T_j(G_M))$. This time, however, since $n \geq 3$, rows $i, j \in [n]$ in $\overline{M}$ both contain at least two distinct elements and thus in $G_{\overline{M}}$, $i, j \in [n]$ are connected to $1 \in \mu_n$ as well as another vertex in μ_n. We reset our notation so M is the matrix with the necessary switches such that G_M has our desired properties.

Since G_M satisfies these properties, if there is a vertex in $[n]$ which is only connected to $1 \in \mu_n$, denote it v, otherwise let v be any vertex connected to $1 \in \mu_n$. We now define the graph $G'_M \coloneqq G_M[([n] \backslash \{v\}, \mu_n \backslash \{1\})]$. The reason we performed all these tedious switches is so we can guarantee that every vertex in $[n] \backslash \{v\} \subset V(G'_M)$ has at least one edge, implying none of the rows are made up of only one distinct element.

If there is a cluster in G'_M, there is a minimal cluster in G'_M from which we pick a vertex i. We know that none of the rows of M are made up of only one distinct element, including row i, so by Corollary 2.37, we can swap two distinct adjacent elements, giving us a new matrix $\tilde{M}$ where $G'_{\tilde{M}} \coloneqq G_{\tilde{M}}[([n]\backslash\{v\}, \mu_n\backslash\{1\})] \in T_i(G'_M)$. Also, by Lemma 2.32, $\mathrm{cd}(G'_{\tilde{M}}; [n]\backslash\{v\}) < \mathrm{cd}(G'_M; [n]\backslash\{v\})$. Repeating the above process, resetting our notation back to M every time, we slowly untangle clusters in the graph G'_M thus reducing $\mathrm{cd}(G'_M; [n]\backslash\{v\})$ until $\mathrm{cd}(G'_M; [n]\backslash\{v\}) = 0$ and thus by Remark 2.29, we have a perfect matching on G'_M. At this point, we will see that we now have a perfect matching on G_M, since $(v, 1) \in E(G_M)$. Now, by Theorem 2.16, since $\mathrm{char}(\mathbb{F}) \nmid n$, M can be unlocked by row rotations and after applying these rotations, our n^2 elements make up an invertible matrix. □

Remark 2.38. *It is very easy to apply the proof of the above Theorem, Theorem 2.30, to the $n = 2$ case since we only use the condition that $n \geq 3$ once. If we rewrite the third paragraph of the proof in the $n = 2$ case, it is easy to see that the only case we need to reconsider is when we have matrix $M = \begin{pmatrix} a & -a \\ -a & a \end{pmatrix}$ for any $a \in \mathbb{F}$ and thus $E(G_M) = \{(1, -1), (2, -1)\}$. In this case, when we swap two different elements, we get $\hat{M} = \begin{pmatrix} a & -a \\ a & -a \end{pmatrix}$. Now $E(G_{\hat{M}}) = \{(1, 1), (2, 1)\}$ so $G_{\hat{M}}$ now has two vertices connected to $1 \in \mu_2$ and by the proof, we should switch a and $-a$ back. However, when we switch the other a and $-a$, we're back to the matrix M and thus caught in an infinite loop. In fact, it is no wonder the proof doesn't work for this matrix, since we can never construct an invertible matrix if we're given the elements $\{a, a, -a, -a\}$! So, by considering this case separately, we can actually give the following statement.*

For $n \in \mathbb{N}$, given n^2 elements in a field $\mathbb{F}$, then, unless there are more than $n^2 - n + 1$ of the same element, more than $n^2 - n$ zeroes or the elements are $\{a, a, -a, -a\}$ for some $a \in \mathbb{F}$, we can always construct an invertible $n \times n$ matrix out of those elements.

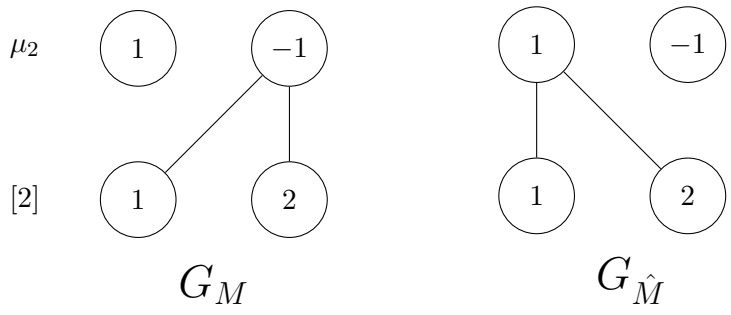

Figure 2: The bipartite graphs G_M and $G_{\hat{M}}$ corresponding to M and $\hat{M}$ respectively from Remark 2.38.

Corollary 2.39. *For $n \in \mathbb{N}$, given n^2 elements in a field $\mathbb{F}$, we can always construct an invertible $n \times n$ matrix out of those elements if and only if there are at most $n^2 - n + 1$ of the same element, at most $n^2 - n$ zeroes and the elements are not $\{a, a, -a, -a\}$ for some $a \in \mathbb{F}$.*

Proof. Then $n = 1$ case is trivial. For $n \geq 2$, the $\Leftarrow$ follows from Theorem 2.30, Theorem 2.19 and Remark 2.38. For the $\Rightarrow$ direction, if there are more than $n^2 - n + 1$ of the same element, by the pigeonhole principle, there are always going to be two rows filled with only one distinct element no matter how we rearrange the matrix. Thus, since those two rows are not linearly independent, the determinant of the matrix will always vanish. Similarly, if there are more than $n^2 - n$ zeroes, again by the pigeonhole principle, there will be at least one row made up of just zeroes and thus the determinant

will always be zero. Finally, for $a \in \mathbb{F}$, the determinants of $\begin{pmatrix} a & a \\ -a & -a \end{pmatrix}$, $\begin{pmatrix} a & -a \\ -a & a \end{pmatrix}$ and $\begin{pmatrix} a & -a \\ a & -a \end{pmatrix}$ are all zero so if the elements are $\{a, a, -a, -a\}$, we cannot rearrange the elements such that the matrix is invertible. □

Remark 2.40. *Although we are talking about being able to construct invertible $n \times n$ matrices from n^2 elements in Theorem 2.19, Theorem 2.30 and Corollary 2.39, this is equivalent to saying matrices made up of those n^2 elements can be unlocked by all permutations.*

Remark 2.41. *Given n^2 elements in a field $\mathbb{F}$ which can be arranged into an invertible $n \times n$ matrix as in Corollary 2.39, the proofs of Theorem 2.30 and Theorem 2.19 in fact both give algorithms to find an invertible $n \times n$ matrix constructed from those elements when combined with Remark 2.8.*

2.3 ...by rotations of its rows and columns?

We now consider rotations of both rows and columns of our matrix. Letting e_i denote the standard ith basis vector as before, let $c_i(M) \coloneqq (r_i(M^T))^T$ ie. a rotation of the ith column by 1 element. Then $c_i \in S_{n^2}$ and for $C \coloneqq \{c_i : i \in [n]\} \subseteq S_{n^2}$, then $\langle R, C \rangle \subseteq S_{n^2}$ is the set of all row and column rotations. We say M is unlocked by row and column rotations if $\exists\, \sigma \in \langle R, C \rangle$ such that $\det(\sigma(M)) \neq 0$.

Example 2.42. *Let's consider the matrix $\gamma = \begin{pmatrix} 2 & -7 & 5 \\ -3 & -8 & 11 \\ 2 & 0 & -2 \end{pmatrix} \in M_3(\mathbb{C})$ again. From Example 2.3 we know we won't get a non-zero determinant by rotating the rows, however, by rotating the first column, we get $c_1(\gamma) = \begin{pmatrix} -3 & -7 & 5 \\ 2 & -8 & 11 \\ 2 & 0 & -2 \end{pmatrix}$ which $\det(c_1(\gamma)) = -150$. Now consider $\nu = \begin{pmatrix} -8 & 1 & 7 \\ 6 & 4 & 9 \\ 2 & -5 & 3 \end{pmatrix} \in M_3(\mathbb{F}_{19})$. It is easy to see that ν is not unlocked by just rows or just columns ie. $\det(\sigma(\nu)) = 0$ for all $\sigma \in \langle R \rangle \cup \langle C \rangle$ since both the rows and columns add up to 0. However, if we rotate the first column down by one we have $c_1^2(\nu) = \begin{pmatrix} 2 & 1 & 7 \\ -8 & 4 & 9 \\ 6 & -5 & 3 \end{pmatrix}$ and then if we rotate the top row by one we get $r_1^2 c_1^2(\nu) = \begin{pmatrix} 7 & 2 & 1 \\ -8 & 4 & 9 \\ 6 & -5 & 3 \end{pmatrix}$ which has determinant $1 \neq 0$.*

We present a final original theorem, building on our work to which matrices are unlocked by all permutations.

Theorem 2.43. *For $n \geq 3$, given a matrix $M \in M_n(\mathbb{F})$, then M is unlocked by row and column rotations if and only if there are at most $n^2 - n + 1$ of the same element or at most $n^2 - n$ zeroes.*

As we can see, by comparing with Theorem 2.30, allowing rotations of both row and columns allows us as much freedom as rearranging the elements in any way we like. To see why, we need to briefly revisit some group theory of the symmetric groups S_n.

Lemma 2.44. *By rotating the rows and columns of a matrix $M \in M_n(\mathbb{F})$, we can cyclically permute any n elements of the matrix in any order leaving all other entries unchanged.*

Proof. Choose n elements in the matrix that we want to cyclically permute. We now want to manoeuvre all of these elements into the first row in the specified order just by using rotations of rows and columns. In the specified order that we want our elements to be cyclically permuted, we move one element at a

time, making sure not to alter any of the elements already correctly placed in the first row. Say we have an element currently at (i, j) which we want to move to position $(1, k)$. If $i = 1$, then rotate the jth column by one position and let the element be at (i, j) where $i \neq 1$. Now rotate the ith row until the element is in column k and rotate the kth column until the element is in $(1, k)$. In this process, we do not alter the position of any element already placed correctly in the first row. By repeating this for all n elements we want to cyclically permute, we have a sequence of rotations which get us from our starting matrix M to the matrix with the elements we want to cyclically permute in the first row in the specified order. Now, we rotate the top row to cyclically permute the elements as specified and then perform the inverse of each rotation in our sequence in reverse order to get back to the starting matrix M now with those n elements cyclically permuted. Clearly this was done only using rotation of rows and columns since the inverse of a rotation is again a rotation. □

We will now need a standard result from group theory which we will not prove here, however, a proof can be found in Cook's lecture notes, [Coo10].

Lemma 2.45. *For $n \geq 5$, the only normal subgroups of S_n are $\{e\}, A_n$ and S_n.*

We now prove another group theory result utilising Lemma 2.45.

Lemma 2.46. *For $n \geq 3$,*

$$n\text{-cycles in } S_{n^2} \text{ generate } \begin{cases} A_{n^2} & \text{if } n \text{ is odd,} \\ S_{n^2} & \text{if } n \text{ is even.} \end{cases}$$

Proof. Since conjugacy classes in S_{n^2} are given by elements with the same cycle shape, the set of $n-$cycles form a full conjugacy class. Letting the subgroup generated by the $n-$cycles be denoted N, then N is normal in S_{n^2}. To see this, notice that for all $g \in S_{n^2}$,

$$g^{-1}Ng = g^{-1}\left\{\prod a : a \text{ is an } n-\text{cycle}\right\} g = \left\{\prod g^{-1}ag : a \text{ is an } n-\text{cycle}\right\} \subseteq N$$

since the $n-$cycles form a conjugacy class and thus $g^{-1}ag$ is again an $n-$cycle for any $g \in S_{n^2}$.

Now we use Lemma 2.45 and notice that for n odd, an $n-$cycle is an even permutation ie. has $\text{sgn}(\sigma) = 1$, and since sgn is multiplicative, N will only contain elements $\sigma \in S_{n^2}$ with $\text{sgn}(\sigma) = 1$ ie. even permutations. Thus, for n odd, $N \neq S_{n^2}$ as S_{n^2} contains odd permutations and N is clearly not the trivial subgroup then $N = A_{n^2}$. Similarly, for n even, $N \not\subseteq A_{n^2}$ since, for n even, an $n-$cycle is an odd permutation thus $N = S_{n^2}$. □

We will now prove Theorem 2.43 by using the fact that the n-cycles generate either A_{n^2} or S_{n^2} and then using Corollary 2.39.

Proof of Theorem 2.43. By Lemma 2.44, if we think of S_{n^2} acting on each of the elements in M, then the $n-$cycles given by rotations of the rows and columns generate all $n-$cycles in S_{n^2}. Now, combining this with Lemma 2.46, we have that for n even, we can apply any permutation to the elements of M just by rotating the rows and columns and for n odd, we can apply any even permutation to the elements of M just by rotating the rows and columns.

Now, the result follows for n even using Corollary 2.39, and we only have to work slightly harder for n odd. In this case, if we recall that, in the process of proving both Theorem 2.30 and Theorem 2.19, we proved that by swapping elements, we could rearrange M into a matrix with non-zero determinant. If this permutation is even, we are done since this permutation can be realised by rotating rows and

columns. If this permutation is odd, we choose two rows and swap all pairs of elements in the same column between the rows. Since n is odd we are adding an odd number of transpositions to the odd permutation thus leaving us with an even permutation. Swapping two rows of a matrix multiplies the determinant by -1, thus, it stays non-zero, and since we applied an even permutation to get matrix into this form, we could also have got to this point by rotating rows and columns. □

Corollary 2.47. *For $n \in \mathbb{N}$, given a matrix $M \in M_n(\mathbb{F})$, then M is unlocked by row and column rotations if and only if there are at most $n^2 - n + 1$ of the same element, at most $n^2 - n$ zeroes and the elements are not $\{a, a, -a, -a\}$ for some $a \in \mathbb{F}$.*

Proof. Again, the $n = 1$ case is trivial and $n \geq 3$ is given by Theorem 2.43. Extending to the case $n = 2$, the normal subgroups of S_4 are given by $\{e\}, V_4, A_4, S_4$ where $V_4 \subseteq A_4$ so since a $2-$cycle is an odd permutation and using the fact that the rotations of rows and columns generates a normal subgroup, then any permutation of the elements of a 2×2 matrix can be generated by rotations of the rows and columns. The result now follows using Corollary 2.39. □

3 Further Directions for Research

We conclude this paper by discussing 3 further directions that research could be taken in, along the same lines as that of this paper.

- We gave exact conditions on when matrices could be unlocked by $\langle R \rangle$ (row rotations), $\langle R, C \rangle$ (row and column rotations) and S_{n^2} (all permutations) where we circumvented the proof for row and column rotations by proving that $\langle R, C \rangle$ was equal to either A_{n^2} or S_{n^2} depending on the parity of n and then appropriating our proof for all permutations. An obvious route for further research would be to give conditions on when matrices can be unlocked by other subsets of S_{n^2}. We assume that taking subsets such as $\langle R \rangle$ with nice properties will give nicer results.

- In [BKS14], for matrices M, constructed directly from bipartite graphs, a formula is given for the number of row rotations that unlock the matrix, given by $\text{supp}(\hat{f}_M)$ where $\hat{f}_M$ is the discrete Fourier transform of f_M. An investigation into whether a similar formula could be given for any matrix M would likely be successful. However, trying to find formulae for the number of elements of other subsets of S_{n^2} that the matrix is unlocked by seems fruitless based on the slightly arduous proof of even one of these elements existing.

- The polynomial ideal $\mathcal{J}(n)$ defined below is the subject of Kézdy and Snevily's paper titled *Polynomials that Vanish on Distinct Roots of Unity*, [KS04], where amongst other things, they give a Gröbner basis for the ideal $\mathcal{J}(n)$ and use Gröbner basis methods to give a characterisation of $\mathcal{J}(n)$ based on the Combinatorial Nullstellensatz.

  Definition 3.1. *Let $\mathcal{J}(n)$ be an ideal in $\mathbb{C}[x_1, ..., x_n]$, where $g \in \mathcal{J}(n) \Leftrightarrow g(x) = 0$ for all $x \in \mu_n^n$ with distinct components ie. $x_i \neq x_j$ for $i \neq j$.*

  It is easy to see that for $g \in \mathbb{C}[x_1, ..., x_n]$, let $f(x) := g(x)(\det V_n)(x)$, then, $g \in \mathcal{J}(n) \Leftrightarrow \mu_n^n \subseteq Z(f)$ ie. $f(x) = 0$ for all $x \in \mu_n^n$. Combining the above with Lemma 2.13 we get $g \in \mathcal{J}(n) \Leftrightarrow (\det V_n)g \in \langle x_i^n - 1 : i \in [n] \rangle$. This is given as a Remark on p.54 of [KS04] and should make the definition of $f_M(x)$ in Definition 2.4 slightly less arbitrary. We now recall the graph polynomial f_G for some graph G.

Definition 3.2. *The graph polynomial* $f_G \in \mathbb{C}[x_1, ..., x_n]$ *of a graph* $G = (V, E)$ *where we enumerate* $V = \{v_i\}_{i=1}^n$ *is given by*

$$f_G(x) \coloneqq \prod_{\substack{(v_i,v_j)\in E \\ i<j}} (x_j - x_i).$$

Theorem 3.3. *[Alo99] Let* $G = (V, E)$ *be a graph where we enumerate* $V = \{v_i\}_{i=1}^n$*. Then* G *is* k*-colourable if and only if the graph polynomial* $f_G \notin \langle x_i^k - 1 : i \in [n]\rangle$*.*

It is interesting to note that, denoting K_n the complete graph on n nodes and V_n the Vandermonde matrix, $\det V_n(x) = f_{K_n}(x)$. Thus, for a graph G with vertices $[n]$, we could define the ideal $\mathcal{J}_G(n)$ where $g \in \mathcal{J}_G(n) \Leftrightarrow f_G g \in \langle x_i^n - 1 : i \in [n]\rangle$. Then, it is clear that $\mathcal{J}_{K_n}(n) = \mathcal{J}(n)$. $\mathcal{J}_G(n)$ would then have the property that $g \in \mathcal{J}(n) \Leftrightarrow g(x) = 0$ for all $x \in \mu_n^n$ where $x_i \neq x_j$ if $(i, j) \in E(G)$. We can take this idea a step further by extending to hypergraphs. A similar result to Theorem 3.3 which can be found in Alon's paper, [Alo99], gives an exact condition on when an 3-uniform hypergraph is 2-colourable. This result can easily be generalised to any hypergraph, not just 3-uniform ones.

Theorem 3.4. *For* $k, n \in \mathbb{N}$*, let* ω *be a primitive* k*th root of unity in* $\mathbb{C}$ *and let* $H = (V, E)$ *be a hypergraph where we enumerate the vertices* $V = \{v_i\}_{i=1}^n$*. Now define a polynomial* $g_H \in \mathbb{C}[x_1, ..., x_n]$ *where*

$$g_H(x) = \prod_{e\in E} \prod_{\tau\in\mu_k} \left(\left(\sum_{v_i \in e} x_i \right) - |e|\tau \right)$$

Then $H = (V, E)$ *is* k*-colourable if and only if* $g_H \notin \langle x_i^k - 1 : i \in [n]\rangle$*.*

For a hypergraph H with vertices $[n]$, we could define ideals $\mathcal{J}_H(n)$ where $h \in \mathcal{J}(n) \Leftrightarrow g_H h \in \langle x_i^n - 1 : i \in [n]\rangle \Leftrightarrow h(x) = 0$ for all $x \in \mu_n^n$ where $|\{x_i : i \in e\}| \neq 1, \forall\, e \in E(H)$.
Gröbner bases for these ideals could likely be found by hand and could certainly be computed. Perhaps even exact conditions on membership of these ideals could be constructed as was done for $\mathcal{J}(n)$ in [KS04]. One application, similar to those of [KS04], that could be established is the following:
Let $G = (V, E)$ *be the bipartite graph where* $V = ([n], \mu_n)$*, and* ω *is a primitive* n*th root of unity. Then define a polynomial* $g \in \mathbb{C}[x_1, ..., x_n]$ *where* $\prod_{(i,\omega^j)\notin E(G)}(x_i - \omega^j)$*. By Lemma 2.15, we have that* G *has a perfect matching* $\Leftrightarrow g \notin \mathcal{J}(n)$*. Now let* $H_{(n,3)}$ *denote the hypergraph on* n *nodes where* $E(H_{(n,3)})$ *contains every possible hyperedge of size* 3 *and no others. Then* G *contains a matching where we allow at most two nodes in* $[n]$ *to connect to the same node in* $\mu_n \Leftrightarrow g \notin \mathcal{J}_{H_{(n,3)}}(n)$*.*

Bibliography

[Alo99] Noga Alon. "Combinatorial Nullstellensatz". In: *Combinatorics, Probability and Computing* 8.1-2 (1999), pp. 7–29. DOI: 10.1017/S0963548398003411.

[BKS14] Timothy Brauch, André Kézdy, and Hunter Snevily. "The Combinatorial Nullstellensatz and DFT on Perfect Matchings in Bipartite Graphs". In: *Ars Combinatoria* 114 (2014), pp. 461–475.

[Bru92] A. A. Bruen. "Polynomial multiplicities over finite fields and intersection sets". In: *Journal of Combinatorial Theory, Series A* 60.1 (1992), pp. 19–33. DOI: 10.1016/0097-3165(92)90035-S.

[Coo10] Bill Cook. *A_n is simple*. 2010. URL: https://www.billcookmath.com/courses/math4720-fall2010/math4720-fall2010-An_is_simple.pdf.

[DeV] M. DeVos. *Graph Theory*. URL: https://www.sfu.ca/~mdevos/notes/graph/345_matchings.pdf.

[Hal86] Marshall Jr. Hall. "Combinatorial Theory". In: *New York: John Wiley and Sons* (1986).

[Hir07] Jonathan Hirata. *Notes on Matching*. 2007. URL: https://math.mit.edu/~djk/18.310/Lecture-Notes/MatchingProblem.pdf.

[KS04] André Kézdy and Hunter Snevily. "Polynomials that Vanish on Distinct nth Roots of Unity." In: *Combinatorics, Probability and Computing* 13 (2004), pp. 37–59. DOI: 10.1017/S0963548303005923.

[Ore55] Oystein Ore. "Graphs and matching theorems". In: *Duke Mathematical Journal* 22.4 (1955), pp. 625–639. DOI: 10.1215/S0012-7094-55-02268-7.

[TV06] Terence Tao and Van H. Vu. *Additive Combinatorics*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2006. DOI: 10.1017/CBO9780511755149.