20.2. DIMENSION AND POLYNOMIALS

The dimension of a vector space is an upper bound on the number of vectors in any linearly independent set. This fundamental and innocuous statement in linear algebra can be very effective in solving extremal problems in combinatorics. The art in applying this method is in designing an appropriate vector space in which the desired set corresponds to linearly independent vectors.

Often we use a space consisting of polynomials, viewing a polynomial as a linear combination of monomials. The unpublished book Babai–Frankl [1992] provides a thorough introduction to the resulting "polynomial method", along with many other aspects of the use of linear algebra in combinatorics.

THE POLYNOMIAL METHOD

Following Babai–Frankl [1992], we begin with two easy examples. They illustrate both the elegance of dimensionality proofs and the process of turning extremal problems into dimension problems.

20.2.1. Example. *Eventown vs. Oddtown.* A town with n people contains many clubs such that every two clubs have an even number of common members. How many clubs can there be if all the clubs have even size? How many if all the clubs have odd size?

When the clubs have even size ("Eventown"), it is easy to form $2^{\lfloor n/2 \rfloor}$ clubs. Simply group the residents into pairs and form each club as a subset of these pairs. In fact, this is optimal (Exercise 5).

When the clubs have odd size ("Oddtown"), we can form n clubs by using clubs of size 1, or clubs of size $n-1$, or other constructions. In fact, there are between $2^{n(n+2)/8}/(n!)^2$ and $2^{n^2}/n!$ nonisomorphic constructions of size n (Exercise 2). These are much smaller than the Eventown constructions, but nevertheless we show next that they are optimal. ■

The **incidence vector** u of a subset A of $[n]$ is the binary n-tuple recording membership in A: $u_i = 1$ if $i \in A$, otherwise $u_i = 0$. The simple but important observation that allows us to convert problems about intersections of sets into algebraic problems is that if u and v are the incidence vectors of subsets A and B on $[n]$, then $u \cdot v = |A \cap B|$, where $u \cdot v$ is the ordinary **dot product** of u and v: $u \cdot v = \sum_{i=1}^{n} u_i v_i$.

20.2.2. THEOREM. (Berlekamp [1969]) If $\mathcal{F}$ is a family of odd-size subsets of $[n]$ whose pairwise intersections have even size, then $|\mathcal{F}| \le n$.

Proof: Let $\mathcal{F} = \{A_1, \dots, A_m\}$. It suffices to show that the corresponding incidence vectors $u^{(1)}, \dots, u^{(m)}$ are linearly independent, since every n-dimensional vector space has at most n linearly independent vectors. The proof is simplest when we use $\mathbb{F}_2$ as the field over which the space is defined, since we have conditions on the parity of the intersections. Thus all our numerical computations with vectors in $\mathbb{F}_2^n$ are modulo 2.

The conditions on the sizes of the sets and their pairwise intersections require that $u^{(i)} \cdot u^{(i)} \equiv 1$ for $1 \le i \le m$ and $u^{(i)} \cdot u^{(j)} \equiv 0$ for $i \ne j$.

To prove that the vectors are linearly independent, we form an equation of dependence: $\sum_{i=1}^{m} c_i u^{(i)} = \mathbf{0}$. When we take the dot product of both sides with $u^{(j)}$, the conditions on the dot products imply that the equation becomes $c_j u^{(j)} \cdot u^{(j)} = 0$ and then $c_j = 0$. ■

The argument made above using the dot product extends to more general functions. When u is a fixed n-tuple in $\mathbb{F}^n$, we can view $u \cdot x$ as a polynomial function of x; that is, $u \cdot x \in \mathbb{F}[x_1, \dots, x_n]$. We obtain a polynomial in n variables; it has degree 1 in each variable. The general form of the preceding argument is the following.

20.2.3. PROPOSITION. (The **Diagonal Criterion**) Let $f_1, \dots, f_m$ be functions in a linear space. If $v^{(1)}, \dots, v^{(m)}$ are points such that $f_i(v^{(i)}) \ne 0$ for $1 \le i \le m$ and $f_i(v^{(j)}) = 0$ for $i \ne j$, then $f_1, \dots, f_m$ are linearly independent.

Proof: Consider $c_1, \dots, c_m$ such that $\sum_{i=1}^{m} c_i f_i$ is identically zero. Evaluating $\sum_{i=1}^{m} c_i f_i$ at $v^{(j)}$ yields $c_j f_j(v^{(j)}) = 0$ and hence $c_j = 0$. ■

Our next application of the diagonal criterion shows the polynomial method more fully.

20.2.4. DEFINITION. A **k-distance set** is a set of points such that the distances between points lie in a set of at most k numbers.

For example, a one-distance set in $\mathbb{R}^n$ must lie at the vertices of a simplex, so the size of a 1-distance set in $\mathbb{R}^n$ is at most $n-1$. For a two-distance set, Exercise 6 requests a construction for a lower bound of $\binom{n+1}{2}$. We next prove an upper bound that is not much larger.

20.2.5. THEOREM. (Larman–Rogers–Seidel [1977]) Every two-distance set in $\mathbb{R}^n$ has at most $(n+1)(n+4)/2$ points.

Proof: Let $\{v^{(1)},\dots,v^{(m)}\}$ be a two-distance set, and let c and d be the two distances. To avoid square roots, we compute with squared distances. Write $\|x-y\|^2$ for the square of the distance between x and y; it equals $\sum_{j=1}^n (x_j-y_j)^2$.

Define polynomials $f_1,\dots,f_m$ by $f_i(x)=(\|x-v^{(i)}\|^2-c^2)(\|x-v^{(i)}\|^2-d^2)$. Note that $f_i(v^{(i)})=c^2d^2\neq 0$, and $f_i(v^{(j)})=0$ for $i\neq j$, since $\|v^{(j)}-v^{(i)}\|\in\{c,d\}$. By the diagonal criterion, $f_1,\dots,f_m$ are linearly independent.

To obtain a good bound on m, we want to capture $f_1,\dots,f_m$ within a small subspace of $\mathbb{R}[x_1,\dots,x_n]$. Written as a polynomial, we have

$$f_i(x)=\Big(\sum_{k=1}^n (x_k-v_k^{(i)})^2-c^2\Big)\Big(\sum_{k=1}^n (x_k-v_k^{(i)})^2-d^2\Big).$$

When expanded completely, the total degree in each monomial term is at most 4. The polynomial is a linear combination of such monomials. The number of ways to distribute total degree at most 4 over n variables, forming such monomials, is less than n^4. Hence $m<n^4$.

To prove a better bound, we capture f_i in the span of fewer monomials. When expanding the product, the terms with degree 4 are $\big(\sum_{k=1}^n x_k^2\big)^2$. Those with degree 3 have the form $x_j\big(\sum_{k=1}^n x_k^2\big)$. Thus f_i is a linear combination of polynomials of the following forms:

$$\Big(\sum_{k=1}^n x_k^2\Big)^2,\quad x_j\Big(\sum_{k=1}^n x_k^2\Big)^2,\quad x_jx_k,\quad x_j,\quad 1,$$

where $j,k\in[n]$. The number of such polynomials is $1+n+n(n+1)/2+n+1$, which simplifies to $(n+1)(n+4)/2$. ■

20.2.6. REMARK. *The polynomial method.* The proof of Theorem 20.2.5 illustrates the general outline of the polynomial method to show that a set S has size at most m.

1) Define polynomials associated with the elements of S.
2) Show that the polynomials are linearly independent.
3) Show that the polynomials are spanned by a set of size m.

Step 3 shows that the polynomials lie in a space of dimension at most m. Since they are linearly independent, there are at most m of them.

Occasionally one can prove a better bound by adding a step 2.5, which is to throw in additional polynomials besides the ones associated with S and show that the polynomials in the augmented family remain linearly independent. Blokhuis [1981] did this to improve the bound in Theorem 20.2.5 from $(n+1)(n+4)/2$ to $(n+1)(n+2)/2$. In addition to the polynomials $f_1,\dots,f_m$ defined there, he added the constant polynomial 1 and the linear polynomials $x_1,\dots,x_n$ of degree 1. The full set is spanned by the same polynomials as before and is linearly independent, so the bound on m is reduced by $n+1$ (see Exercise 7).

We will use this augmentation technique in Theorem 20.2.@. ■

Often we need an analogous criterion for linear independence that holds more generally than the diagonal criterion.

20.2.7. PROPOSITION. (The **Triangular Criterion**) Let $f_1,\dots,f_m$ be functions in a linear space. If $v^{(1)},\dots,v^{(m)}$ are points such that $f_i(v^{(i)})\neq 0$ for $1\le i\le m$ and $f_i(v^{(j)})=0$ for $i>j$, then $f_1,\dots,f_m$ are linearly independent.

Proof: Consider coefficients $c_1,\dots,c_m$ such that $\sum_{i=1}^m c_if_i$ is the identically-zero function. Evaluating this function at $v^{(1)}$ yields $c_1f_1(v^{(1)})=0$ and hence $c_1=0$. Proceding by induction on j, if we have already verified that $c_1=\cdots=c_{j-1}=0$, then evaluating $\sum_{i=1}^m c_if_i$ at $v^{(j)}$ yields $c_jf_j(v^{(j)})=0$, because the earlier terms have coefficient 0 and the later functions evaluate to 0. Hence $c_j=0$ for $1\le j\le m$. ■

FAMILIES WITH RESTRICTED INTERSECTIONS

Restricting the sizes of the intersections of sets in a family restricts the size of the family. Perhaps the most famous such result is the Erdős–Ko–Rado Theorem [1961]. A family of sets is an **intersecting family** if every two members have a nonempty intersection. For $n\ge 2k$, a consequence of their theorem is that the maximum size of an intersecting family of k-sets in $[n]$ is $\binom{n-1}{k-1}$ (see Chapter 13). More generally, we could specify the allowed sizes of intersections.

20.2.8. DEFINITION. For $L\subseteq\mathbb{N}_0$, an *L-intersecting family* of sets is a family $\mathcal{F}$ such that $|A\cap B|\in L$ for all $A,B\in\mathcal{F}$.

In this language, the Erdős–Ko–Rado Theorem for k-uniform families uses $L=\{1,\dots,k-1\}$. If 0, so that $L=\{0,\dots,k-1\}$, then we can include all k-sets, yielding a family of size $\binom{n}{|L|}$. If we do not require a k-uniform family, then we can include all sets of size at most k, yielding a family of size $\sum_{i=0}^{|L|}\binom{n}{i}$. Frankl–Wilson [1981] proved that for all n and L, no L-intersecting family has more than this many members.

To prove the Frankl–Wilson Theorem, we will use the Triangular Criterion and another method for bounding the size of a set of linearly independent polynomials. We modify them, without changing the values responsible for making them linearly independent, into other polynomials contained in a space of small dimension.

20.2.9. REMARK. *Multilinear Reduction method.* Given a polynomial f in n variables, define the **multilinear reduction** of f to be the polynomial $\hat{f}$ in which each positive exponent (in the expression of f as a sum of monomials) is reduced to 1.

Because $0^r = 0$ and $1^r = 1$ for $r \in \mathbb{N}$, the values of f and $\hat{f}$ agree on $\{0,1\}^n$. If $f_1, \ldots, f_m$ are linearly independent due to their values on $\{0,1\}^n$, then $\hat{f}_1, \ldots, \hat{f}_m$ are linearly independent for the same reason. In general, better bounds are available on the number of linearly independent multilinear polynomials. ■

20.2.10. THEOREM. (Frankl–Wilson [1981]) If $\mathcal{F}$ is an L-intersecting family of subsets of $[n]$, where $|L| = s$, then $|\mathcal{F}| \le \sum_{i=0}^{s}\binom{n}{i}$.

Proof: (Babai [1988]) Let $\mathcal{F} = \{A_1, \ldots, A_m\}$, indexed so that $|A_1| \le \cdots \le |A_m|$. Let $L = \{l_1, \ldots, l_s\}$. For each i, let v_i be the incidence vector of A_i. Define polynomials $f_1, \ldots, f_m$ by $f_i(x) = \prod_{k\colon l_k<|A_i|}(x \cdot v_i - l_k)$. By construction, $f_i(v_j) \ne 0$ for $j = i$. Using the indexing of $\mathcal{F}$ by size, $|A_j \cap A_i| < |A_i|$ for $j < i$, and hence $f_i(v_j) = 0$ for $j < i$. By the Triangular Criterion, $f_1, \ldots, f_m$ are linearly independent.

Since $v_1, \ldots, v_m \in \{0,1\}^n$, the computations hold also for the multilinear reductions $\hat{f}_1, \ldots, \hat{f}_m$, so these polynomials also are linearly independent. We prove the desired bound by capturing *them* in a small space. Because each f_i is the product of s linear factors, the total degree of each monomial in the expansion of f_i is bounded by s. Hence the multilinear reduction of f_i is spanned by the monomials that are products of at most s distinct variables. The number of such monomials is $\sum_{i=0}^{s}\binom{n}{i}$. ■

Better bounds can be proved in special cases. For example, if the sets have odd size and the elements of L are all even, then the Oddtown theorem (Theorem 20.2.2) implies $|\mathcal{F}| \le n$. This is consistent with the bound $\binom{n}{|L|}$ if we view the intersection sizes as congruence classes modulo 2. The class 0 contains all intersection sizes, but the sizes of the sets in the family are forbidden from that class. Exercise 3 generalizes the Oddtown theorem to other moduli for uniform families.

Meanwhile, here we present a modular version of the non-uniform Theorem 20.2.10. The proof is analogous, and the bound is the same.

20.2.11. DEFINITION. Let p be a prime. For $L \subseteq \mathbb{Z}_p$, we say that $t \in L \pmod{p}$ if $t \equiv l \pmod{p}$ for some $l \in L$. A family $\mathcal{F}$ of subsets of $[n]$ is ***p*-modular *L*-intersecting** if $|A| \notin L \pmod{p}$ for $A \in \mathcal{F}$ and $|A \cap B| \in L \pmod{p}$ for distinct $A, B \in \mathcal{F}$.

20.2.12. THEOREM. (Deza–Frankl–Singhi [1983]) Let p be a prime. If $\mathcal{F}$ is a p-modular L-intersecting family of subsets of $[n]$, where $|L| = s$, then $|\mathcal{F}| \le \sum_{i=0}^{s}\binom{n}{i}$.

Proof: (Alon–Babai–Suzuki [1991]) Let $\mathcal{F} = \{A_1, \ldots, A_m\}$, and let v_i be the incidence vector of A_i. Define polynomials $f_1, \ldots, f_m$ by $f_i(x) = \prod_{l\in L}(x \cdot v_i - l)$. We view all computations over $\mathbb{F}_p$ and write "=" instead of "≡". Note that $v_j \cdot v_i = |A_i \cap A_j|$. Thus $f_i(v_j) \ne 0$ for $i = j$ (since $|A_i| \notin L \pmod{p}$), while $f_i(v_j) = 0$ for $i \ne j$ (since $|A_i \cap A_j| \in L \pmod{p}$). By the Diagonal Criterion, $f_1, \ldots, f_m$ are linearly independent.

The bound now follows in the same way as in Theorem 20.2.10. The multilinear reductions $\hat{f}_1, \ldots, \hat{f}_m$ are linearly independent by the same criterion as $f_1, \ldots, f_m$, and they lie in a space of dimension $\sum_{i=0}^{s}\binom{n}{i}$. ■

Although it may seem that the bound $\sum_{i=0}^{s}\binom{n}{i}$ is much larger than $\binom{n}{s}$, actually it is not when s is not too big.

20.2.13. LEMMA. If $n \ge 2s$ and $s = n/r$, then

$$\sum_{i=0}^{s}\binom{n}{i} \le \binom{n}{s}\left(1 + \frac{s}{n-2s+1}\right) < \binom{n}{s}\left(1 + \frac{1}{r-2}\right).$$

In particular, if $s \le n/3$, then $\sum_{i=0}^{s}\binom{n}{i} < 2\binom{n}{s}$.

Proof: By factoring $\binom{n}{s}$ from each term and then enlarging (and extending) the terms to obtain a geometric series,

$$\begin{aligned}
\sum_{i=0}^{s}\binom{n}{i} &= \binom{n}{s}\left(1 + \frac{s}{n-s+1} + \frac{s(s-1)}{(n-s+1)(n-s+2)} + \cdots\right)\\
&\le \binom{n}{s}\left(1 + \frac{s}{n-s+1} + \frac{s^2}{(n-s+1)^2} + \cdots\right)\\
&= \binom{n}{s}\frac{1}{1-\frac{s}{n-s+1}} = \binom{n}{s}\frac{n-s+1}{n-2s+1} = \binom{n}{s}\left(1 + \frac{s}{n-2s+1}\right)\\
&= \binom{n}{s}\left(1 + \frac{n/r}{n-2n/r+1}\right) < \binom{n}{s}\left(1 + \frac{1}{r-2}\right)
\end{aligned}$$

■

These results yield a constructive superpolynomial lower bound for the diagonal Ramsey number $R(t,t)$. It is not as strong as Erdős' nonconstructive exponential lower bound, but the graphs are explicitly defined. The trivial construction in which one color occupies $(t-1)K_{t-1}$ and the other is the complementary complete $(t-1)$-partite graph shows that $R(t,t) > (t-1)^2$. Nagy [1972] increased the explicit lower bound to $\binom{t^3}{3}$ (Exercise 4). Frankl [1977] constructed graphs showing that $R(t,t) > t^{\omega(t)}$ using Δ-systems (sunflowers), where $\omega(t) \to \infty$ as $t \to \infty$. Frankl–Wilson [1981] obtained similar behavior from p-modular L-intersecting families.

20.2.14. THEOREM. (Frankl–Wilson [1981]) Let p be a prime, and choose $n > 2p^2$. Let G be the graph with vertex set $\binom{[n]}{p^2-1}$ defined by $AB \in E(G)$ if and only if $|A \cap B| \not\equiv -1 \pmod{p}$. The graph G has no homogeneous set with more than $2\binom{n}{p-1}$ vertices. As a consequence, $R(t,t) > t^{(1-\varepsilon)\omega(t)}$, where $\omega(t) = \frac{\ln t}{4 \ln\ln t}$.

Proof: If $A_1, \ldots, A_m$ is a clique in G, then it is a p-modular L-intersecting family, where $L = \{0, \ldots, p-2\}$, because $|A_i| = p^2 - 1 \equiv -1 \pmod{p}$, and $A_iA_j \notin E(G)$ when $|A_i \cap A_j| \equiv -1 \pmod{p}$. With $|L| = p-1$, Theorem 20.2.12 yields $m \le \sum_{i=0}^{p-1}\binom{n}{i} < 2\binom{n}{p-1}$. If $A_1, \ldots, A_m$ is an independent set, then $|A_i \cap A_j| \in \{p-1, 2p-1, \ldots, p^2-p-1\}$. Here $p-1$ intersection sizes are allowed, so Theorem 20.2.10 yields $m \le \sum_{i=0}^{p-1}\binom{n}{i} < 2\binom{n}{p-1}$.

Fixing t, let p be the largest prime such that $2\binom{p^3}{p-1} < t$, and let $n = p^3$. We have shown that $R(t,t) > \binom{n}{p^2-1}$. The choice of p yields $p \sim \frac{\ln t}{2\ln\ln t}$, and then $\binom{p^3}{p^2-1} > t^{(1-\varepsilon)\omega(t)}$, where $\omega(t) = \frac{\ln t}{4\ln\ln t}$. That is, we are comparing roughly $\binom{p^3}{p}$ for t with $\binom{p^3}{p^2}$ for the lower bound on $R(t,t)$. The logarithm of the latter is roughly $p/2$ times the logarithm of the former, so roughly $R(t,t) > t^{p/2}$ (Exercise 10 requests further computational details). ■

Next we present an application of p-modular L-intersecting families to coloring the unit-distance graph in n-dimensional space. The famous Hadwiger–Nelson problem (Hadwiger [1944]) asks for the minimum number of colors needed to label the points of $\mathbb{R}^n$ so that no two points at distance 1 have the same color. For $n = 2$, it has long been known that the answer is in $\{4,5,6,7\}$ (Exercise 11). In general, there is an easy upper bound of $n^{n/2}$ (Exercise 12). Larman–Rogers [1972] presented a quadratic lower bound and an upper bound of $(2\sqrt{2} + o(1))^n$ and conjectured an exponential lower bound. Frankl–Wilson [1981] proved this; we obtain it from a simple corollary of Theorem 20.2.12.

20.2.15. COROLLARY. Let p be a prime, and let $\mathcal{F}$ be a $(2p-1)$-uniform family of subsets of $[4p-1]$. If no two members of $\mathcal{F}$ have exactly $p-1$ common elements, then $|\mathcal{F}| \le 2\binom{4p-1}{p-1} < 1.7548^{4p-1}$.

Proof: Let $L = \{0, \ldots, p-2\}$. The family $\mathcal{F}$ is p-modular L-intersecting since $\mathcal{F}$ is $(2p-1)$-uniform with $2p-1 \notin L \pmod{p}$, and the remaining possible intersection sizes $\{p, \ldots, 2p-2\}$ are congruent to elements of L. Since $|L| = p-1$, Theorem 20.2.12 and Lemma 20.2.13 yield the bound.

Note that $2\binom{4p-1}{p-1} = \frac{1}{2}\binom{4p}{p}$. Using Stirling's Formula (Theorem 16.@.@), $\binom{4p}{p}$ is given approximately by $c(4/3^{3/4})^{4p}/\sqrt{p}$ for some constant c, and hence it is bounded by 1.7548^{4p-1} (since $4/3^{3/4} < 1.7548$). ■

20.2.16. THEOREM. (Frankl–Wilson [1981]) For large n, the chromatic number of the unit-distance graph in $\mathbb{R}^n$ is greater than 1.1397^n.

Proof: Note first that the graph defined using distance d is isomorphic to the unit-distance graph. Hence it suffices to prove the claimed lower bound for a subgraph of the distance-d graph. We use an appropriate d and a vertex set that is a subset of the unit cube.

The squared distance between two points in $\{0,1\}^n$ is the number of coordinates where they differ. Viewed as incidence vectors of subsets A and B of n, the number of coordinates whether they differ is $|A \triangle B|$. If A and B have size k, then $|A \triangle B| = 2(k - |A \cap B|)$. Hence forbidding one distance between the points is equivalent to forbidding one intersection size for the sets. If $k = 2p-1$, then forbidding intersection size $p-1$ is equivalent to forbidding squared distance $2p$.

Let p be the largest prime such that $4p-1 \le n$; we use only $4p-1$ of the coordinates. Let $d = \sqrt{2p}$. By Corollary 20.2.15, the maximum size of an independent set in the subgraph of the distance-d graph induced by the incidence vectors of the $(2p-1)$-sets in $[4p-1]$ is at most $2\binom{4p-1}{p-1}$. Hence the chromatic number is at least $\binom{4p-1}{2p-1}/2\binom{4p-1}{p-1}$. Now $3^{3/4}/2 > 1.1397$ completes the proof, since for $m = n/4$ there is a prime between m and $m - m^{7/12}$ when m is sufficiently large (Huxley [1973]). ■

The lower bound can be improved to about $(1.2)^n$ by choosing p to optimize the ratio $\binom{n}{2p-1}/\binom{n}{p-1}$ (Exercise 9).

Frankl–Füredi [1981] conjectured that the bound in the Frankl–Wilson Theorem can be improved when $L = [s]$, limiting the size of an L-intersecting family to $\sum_{i=0}^{s}\binom{n-1}{s}$ instead of $\sum_{i=0}^{s}\binom{n}{s}$. Ramanan [1997] proved this conjecture. Snevily conjectured that the same bound holds when L is any set of s positive numbers. He proved this for sufficiently large n (Snevily [1994]) and when L is an interval (Snevily [1999]) before proving the full conjecture (Snevily [2003]). Snevily's Theorem easily implies the Frankl–Wilson Theorem (Exercise 13).

The case $s = 1$ was proved by Majumdar [1953] and amounts to the non-uniform Fisher inequality (Exercise 19.1.@), which states that if $\mathcal{B}$ is a family of n blocks in $[v]$ (not necessarily of uniform size), no block equals $[v]$, and every two elements appear in λ common blocks, then $n \ge v$. The dual of this (by transposing the incidence matrix) is the statement that the size of an L-intersecting family of subsets of $[n]$ is at most n when $L = \{\lambda\}$. Since $n = \binom{n-1}{0} + \binom{n-1}{1}$, we have the case $s = 1$.

The details of Snevily's Theorem are a bit long, so we present only a modular version that he proved earlier. The proof illustrates a way to improve bounds from the polynomial method. We start with the same polynomials in the same space as before and add more polynomials spanned by the same set. If the full set of polynomials is still linearly independent, then the new bound on $\mathcal{F}$ is the original bound minus the number of added polynomials. In the present application, our space has dimension $\sum_{i=0}^{s}\binom{n}{s}$ and we add $\sum_{i=1}^{s}\binom{n-1}{i-1}$ polynomials, leaving dimension only $\sum_{i=0}^{s}\binom{n-1}{s}$ for those corresponding to $\mathcal{F}$. This technique is also used in Exercise 7 to improve the bound on two-distance sets.

In the case where the sizes of members of $\mathcal{F}$ do not lie in L, we obtain the non-modular statement of Snevily's Theorem by taking p sufficiently large. Extensions to k-wise intersections appear in Grolmusz–Sudakov [2002] and Cao–Hwang–West [2007]. The next lemma can be strengthened, but this statement suffices for our purposes.

20.2.17. LEMMA. Let $C_1, \ldots, C_t$ be subsets of $[n]$, indexed in nondecreasing size order. If polynomials $h_1, \ldots, h_t$ are defined on $\mathbb{R}^n$ by $h_j(x) = \prod_{r\in C_j} x_j$, then $h_1, \ldots, h_t$ are linearly independent on $\{0,1\}^n$.

Proof: Let w_j be the incidence vector of C_j, so $h_j(w_j) = 1$. If $i > j$, then the indexing of $C_1, \ldots, C_t$ guarantees an element $r \in C_i - C_j$. Now h_i has x_r as a factor, but the value in coordinate r of w_j is 0, so $h_i(w_j) = 0$. By the Triangular Criterion, $\{h_1, \ldots, h_t\}$ is linearly independent. ■

20.2.18. THEOREM. (Snevily [1994]) If $\mathcal{F}$ is a p-modular L-intersecting family of subsets of $[n]$, with $s = |L|$, then $|\mathcal{F}| \le \sum_{i=0}^{s}\binom{n-1}{i}$.

Proof: Let $\mathcal{F} = A_1, \ldots, A_m$, indexed so that $A_1, \ldots, A_q$ omit the element 1 and $A_{q+1}, \ldots, A_m$ contain it. Begin the proof as in Theorem 20.2.12, letting $f_i(x) = \prod_{l\in L}(x\cdot v_i - l)$, where v_i is the incidence vector of A_i. As before, the Diagonal Criterion makes $f_1, \ldots, f_m$ linearly independent. As before, the multilinear reductions $\hat{f}_1, \ldots, \hat{f}_m$ are also linearly independent and are spanned by the $\sum_{i=0}^{s}\binom{n}{i}$ multilinear monomials of degree at most s.

Let $C_1, \ldots, C_t$ be the sets of size less than s in $[n]$ lacking element 1, indexed so $|C_1| \le \cdots \le |C_t|$. Define h_j and g_j by $h_j(x) = \prod_{r\in C_j} x_j$ and $g_j(x) = (x_1 - 1)h_j(x)$. Note that g_j has degree at most s (x_1 appears with degree 2 when $j > r$). The multilinear reduction of g_j is spanned by the same set of $\sum_{i=0}^{s}\binom{n}{i}$ monomials as f_i. Since $t = \sum_{i=1}^{s}\binom{n-1}{i-1}$, it suffices to show that $\{f_1, \ldots, f_m\} \cup \{g_1, \ldots, g_t\}$ is linearly independent.

Let $P = \sum_{i=1}^{m}\alpha_i f_i + \sum_{j=1}^{t}\beta_j g_j$, with each α_i and β_j in $\mathbb{F}_p$. Suppose that P is identically 0. Let $A_i' = A_i \cup \{1\}$, and let y_i be the incidence vector of A_i', for $1 \le i \le m$. Each y_i has 1 in the first coordinate, so the contribution of the second sum to $P(y_i)$ is always 0.

Note that $A_j' \cap A_i = A_j \cap A_i$ if $i \le j$. This holds because $1 \in A_j$ if $j > r$ and $1 \notin A_i$ if $i \le r$. Thus $f_i(y_j) = f_i(v_j)$ for $i \le j$. Since $f_i(v_j) = 0$ when $i \ne j$, evaluating P at $y_m, \ldots, y_1$ successively shows that $\alpha_m, \ldots, \alpha_1 = 0$.

By Lemma 20.2.17, $h_1, \ldots, h_t$ are linearly independent. Multiplying all by $x_1 - 1$ leaves $g_1, \ldots, g_t$ independent. Since $\alpha_m, \ldots, \alpha_1 = 0$, making P identically 0 thus also requires $\beta_1, \ldots, \beta_t = 0$. Hence there is no equation of linear dependence for $\{f_1, \ldots, f_m\} \cup \{g_1, \ldots, g_t\}$. ■

The top two entries in $\sum_{i=0}^{s}\binom{n-1}{i}$ sum to $\binom{n}{s}$. When $\mathcal{F}$ is required to be uniform, this bound suffices, even if 0 is allowed in L. This next result appeared first among those we present on L-intersecting families for $|L| > 1$ (proved in 1969 but not published until 1975), but the use of linear algebra in the original proof was different from the approach we have developed. Combining the ideas we have presented led to a shorter proof.

Like the Frankl–Wilson Theorem, this result can be motivated using the case $L = \{0, \ldots, s-1\}$. If $k = s$, then the condition of being L-intersecting places no restriction on the sets chosen for a k-uniform family, so all $\binom{n}{s}$ sets of size k can be chosen. It is perhaps surprising that the same bound is valid for L-intersecting k-uniform families when L is any set of s nonnegative numbers.

20.2.19. THEOREM. (Ray-Chaudhuri–Wilson [1975]) If $n \ge 2s$, and L is a set of s nonnegative integers, then every L-intersecting k-uniform family of subsets of $[n]$ has size at most $\binom{n}{s}$.

Proof: (Alon–Babai–Suzuki [1991]) We may assume $k \notin L$. Let $\mathcal{F} = \{A_1, \ldots, A_m\}$; all are k-sets, with incidence vectors $v_1, \ldots, v_m$. Let $f_i(x) = \prod_{l\in L}(x\cdot v_i - l)$. By the Diagonal Criterion, $f_1, \ldots, f_m$ are linearly independent on $\{0,1\}^n$. Each f_i has degree s; the multilinear reduction yields polynomials $\hat{f}_1, \ldots, \hat{f}_m$ that are linearly independent on $\{0,1\}^n$ and spanned by the $\sum_{i=0}^{s}\binom{n}{i}$ multilinear monomials with degree at most s.

As in Theorem 20.2.18, we add polynomials to this set. Let $C_1, \dots, C_t$ be the sets of size less than s in $[n]$, indexed in increasing order of size. Define h_j by $h_j = \prod_{r\in C_j} x_r$. By Lemma 20.2.17, $h_1, \dots, h_t$ are linearly independent over $\{0,1\}^n$. Define g_j by $g_j(x) = (x \cdot 1_n - k)h_j(x)$. As in Theorem 20.2.18, we have multiplied independent polynomials by one linear factor, and the resulting polynomials are independent.

Let $P = \sum_{i=1}^{m} \alpha_i f_i + \sum_{j=1}^{t} \beta_j g_j$. Consider coefficients such that P is identically 0. Since $\mathcal{F}$ is k-uniform, the contribution from the second sum is 0 when evaluated at v_i. Since $f_i(v_j) = 0$ when $j \neq i$, evaluating P at v_i thus yields $\alpha_i = 0$. With each α_i being 0, linear independence of $g_1, \dots, g_t$ implies also that each β_j is 0.

We conclude that $\{f_1, \dots, f_m\} \cup \{g_1, \dots, g_t\}$ is linearly independent, where again we take the multilinear reduction of g_j. The degree of g_j is at most s, so these vectors also lie in the span of the $\sum_{i=0}^{s} \binom{n}{i}$. Since $t = \sum_{i=0}^{s-1} \binom{n}{i}$, we conclude that $m \le \binom{n}{s}$. ■

As we have noted, the bound $\binom{n}{s}$ holds with equality when $L = \{0, \dots, s-1\}$ and $k = s$. Is it still achievable or nearly achievable for larger k? Because $\binom{n}{s} \sim n^s/s!$, the construction below is not so much smaller than the upper bound when s and t are fixed and n is large.

20.2.20. THEOREM. For $n \ge 2k^2 \ge 2s^2$ and $L = \{0, \dots, s-1\}$, some k-uniform L-intersecting family $\mathcal{F}$ satisfies $|\mathcal{F}| > (n/2k)^s$.

Proof: Let p be the largest prime bounded by n/k, so $n/2k < p \le n/k$. Fix a k-set A contained in $\mathbb{F}_p$. Since $kp \le n$, we may choose X to be an n-set containing $A \times \mathbb{F}_p$; we will ignore the elements of X outside $A \times \mathbb{F}_p$.

Our family $\mathcal{F}$ will consist of p^s k-sets contained in $A \times \mathbb{F}_p$. Given a polynomial f of degree less than s, let $A_f = \{(a, f(a)) \colon a \in A\}$. There are p^s polynomials over $\mathbb{F}_p$ with degree less than s, so this defines p^s sets of size k. Distinct polynomials of degree d over $\mathbb{F}_p$ agree on at most d points in $\mathbb{F}_p$. Since $k \ge s$, we conclude that the p^s sets constructed are distinct, and any two of them have fewer than s common elements. ■

COMBINATORIAL NULLSTELLENSATZ

The Combinatorial Nullstellensatz is a statement about zeros of multivariable polynomials over a field. Fairly easy to prove, it has found wide-ranging applications in additive number theory, discrete geometry, and various parts of graph theory. The theorem was articulated by Noga Alon and presented at a conference in 1995, although the proceedings with the resulting survey paper was not published until 1999. Nevertheless, Alon had already applied the theorem in at least five different papers with eight different coauthors between 1984 and 1996, proving new results and giving short proofs of old results. Since the publication of the theorem, many other researchers have also employed it.

We need a lemma that is a straightforward inductive generalization to n variables of the familiar statement that a nonzero polynomial of degree d in one variable takes the value 0 at most d times. This itself is proved by induction on d, using the Euclidean algorithm to factor out $x - \alpha$ when α is a root. The discussion is valid when the computations are done in any field, and we simply compute with equalities rather than using congruence notation for finite fields.

20.2.21. LEMMA. Let f be a polynomial in n variables $x_1, \dots, x_n$, over a field $\mathbb{F}$. For each i, let the degree of f as a polynomial in x_i be at most d_i, and let S_i be a set of $d_i + 1$ distinct values in $\mathbb{F}$. If $f(x_1, \dots, x_n) = 0$ for $(x_1, \dots, x_n) \in \prod_{i=1}^{n} S_i$, then f is identically 0.

Proof: We take the result in one variable as the basis for induction on n. For $n > 1$, we collect terms to write f as a polynomial in x_n. That is, $f = \sum_{j=0}^{d_n} f_j(x_1, \dots, x_{n-1}) x_n^j$, where each f_j is a polynomial having degree at most d_i in each x_i. For $(x_1, \dots, x_{n-1}) \in \prod_{i=1}^{n-1} S_i$, evaluating $f_0, \dots, f_{d_n}$ yields a one-variable polynomial in x_n of degree at most d_n. Furthermore, the hypothesis implies that this polynomial is 0 for $x_n \in S_n$.

By the basis step ($n = 1$), the one-variable polynomial we obtain for a fixed $(x_1, \dots, x_{n-1}) \in \prod_{i=1}^{n-1} S_i$ is the zero polynomial. Thus each f_i is 0 at all values in $\prod_{i=1}^{n-1} S_i$. By the induction hypothesis, each f_i is identically zero. Thus the coefficients of f are all zero, and f is identically zero. ■

We would like to conclude that if the coefficient of a term $\prod x_i^{t_i}$ is nonzero in a polynomial f of degree $\sum t_i$, and $|S_i| > t_i$ for all i, then the polynomial is nonzero at some point in $\prod S_i$. However, the lemma does not say this, because other terms may have degree larger than t_i in x_i, for some i. Fortunately, it is not hard to overcome this technicality.

The **degree** of a polynomial is the maximum, over all monomials, of the sum of the exponents on the variables. It is convenient to obtain the coefficient of a monomial $\prod_{i=1}^{n} x_i^{t_i}$ in a polynomial $f(x_1, \dots, x_n)$ using the **coefficient operator** $\left[\prod_{i=1}^{n} x_i^{t_i}\right]$, which for formal power series in one variable was used extensively in Chapter 16.

20.2.22. THEOREM. (**Combinatorial Nullstellensatz**; Alon [1999]) If $\prod_{i=1}^{n} x_i^{t_i}$ is a monomial with nonzero coefficient in a polynomial f

having degree $\sum_{i=1}^{n} t_i$ over a field $\mathbb{F}$, and $S_1, \ldots, S_n$ are sets with $|S_i| > t_i$ for $1 \le i \le n$, then $f(x) \ne 0$ for some $x \in \prod S_i$.

Proof: It suffices to prove the statement when $|S_i| = t_i + 1$ for each i. The idea is to change f into another polynomial $\hat{f}$ that agrees with f on $\prod S_i$ but has degree at most t_i as a polynomial in x_i, for each i. Lemma 20.2.21 then implies that $\hat{f}(x) \ne 0$ for some $x \in \prod S_i$. Since $\hat{f}$ agrees with f on $\prod S_i$, also $f(x) \ne 0$.

For each index i, define a polynomial g_i by $g_i(x) = \prod_{s \in S_i}(x_i - s)$; note that g_i depends only on x_i. It has degree $t_i + 1$ in x_i and degree 0 in other variables. Expanding the product yields $g_i(x) = x_i^{t_i+1} - h_i(x)$, where h_i is a polynomial with degree at most t_i in x_i and degree 0 in other variables.

By definition, $g_i(x) = 0$ for $x \in \prod S_i$, since $x_i \in S_i$ in that case. Therefore, $x_i^{t_i+1} = h_i(x)$ for all $x \in \prod S_i$. This allows us to replace each appearance of a variable having too large an exponent with a polynomial of smaller degree in that variable. By making such a replacement as long as the polynomial still has degree greater than t_i in some x_i, we obtain $\hat{f}$ having degree at most t_i in x_i for each i.

We must also check that $\left[\prod x_i^{t_i}\right]\hat{f}(x) \ne 0$. Since no exponent is too large, we made no change to that term. Also we did not introduce any terms that could cancel it; since f has degree $\sum t_i$, any monomial containing a variable with too large an exponent has some other x_j with exponent less than t_j. Since the substitutions increase no exponents, no substitution can introduce a contribution to $\left[\prod x_i^{t_i}\right]$. ■

One of the first applications was a short proof of the Cauchy–Davenport Theorem of additive number theory. The theorem was originally proved by Cauchy in 1813 and Davenport in 1935. Let A and B be subsets of $\mathbb{Z}_n$, with $|A| = a$ and $|B| = b$. The question is how many elements of $\mathbb{Z}_n$ arise as $x + y$ with $x \in A$ and $y \in B$.

Setting $A = \{0, \ldots, a-1\}$ and $B = \{0, \ldots, b-1\}$ shows that the sum can be as small as $\min\{n, a+b-1\}$. On the other hand, if $a + b > n$, then for any $c \in \mathbb{Z}_n$ the sets A and $\{c - y\colon y \in B\}$ must intersect, and when an element x is in the intersection we have $x \in A$ and $y \in B$ such that $c = x + y$. Hence the number of sums always equals n if $a + b > n$. Finally, note that when $n = 2k$, taking the "even" classes for both A and B yields only even classes as sums, so here the sum can be as small as $n/2$ even though $a = b = n/2$.

This suggests restricting our attention to prime moduli and proving the next theorem.

20.2.23. THEOREM. (**Cauchy–Davenport Theorem**) If p is prime, and $A, B \subseteq \mathbb{Z}_p$ with $|A| = a$ and $|B| = b$ and $a + b \le p$, then the smallest possible size of $\{x + y\colon x \in A, y \in B\}$ is $a + b - 1$.

Proof: By the observation above, it suffices to prove the lower bound. Suppose that there are fewer than $a + b - 1$ sums. Let C be a set of size $a + b - 2$ in $\mathbb{Z}_p$ that contains all the sums. Let $f(x, y) = \prod_{c \in C}(x + y - c)$, over $\mathbb{Z}_p$. We have a polynomial in two variables, and its degree is $a + b - 2$.

We claim that $[x^{a-1}y^{b-1}]f(x, y) = \binom{a+b-2}{a-1} \not\equiv 0 \pmod{p}$. Contributions to this coefficient use x or y in each factor when expanding f, choosing x exactly $a - 1$ times and y exactly $b - 1$ times. The number of ways to do that, each contributing $+1$ to the coefficient, is $\binom{a+b-2}{a-1}$. Finally, that binomial coefficient is nonzero modulo p since $a + b - 2 < p$; there is no factor of p in the numerator and no other way to introduce a factor of p.

Since $|A| = a$ and $|B| = b$, the Combinatorial Nullstellensatz yields $x \in A$ and $y \in B$ such that $f(x, y) \ne 0$. This is a contradiction, since f was constructed to be 0 at all such pairs (x, y). ■

This short proof illustrates the method for applying the Combinatorial Nullstellensatz. Using the set of sums, we design f that is 0 at (x, y) when $x \in A$ and $y \in B$. If the set of sums is too small, then $A \times B$ is too big for f to be identically 0 there when the appropriate coefficient is nonzero.

When $A = B$, the lower bound in Theorem 20.2.23 is $\min\{2\,|A| - 1, p\}$. Erdős and Heilbronn [1964] conjectured that almost as much is forced even without considering contributions of the form $a + a$. Given the ease of proving this from the Combinatorial Nullstellensatz, it is remarkable that the problem was open for 30 years. The original proof used exterior algebra and representation theory of the symmetric group.

20.2.24. THEOREM. (**Erdős–Heilbronn Conjecture**; Dias da Silva–Hamidoune [1994]) If $A \subseteq \mathbb{Z}_p$, where p is prime, and C is the set of sums of distinct elements of A, then $|C| \ge \min\{2\,|A| - 3, p\}$.

Proof: (Alon–Nathanson–Rusza [1996]) Since there are only p classes, we may assume that $2a - 3 < p$, where $a = |A|$. As in the proof of Theorem 20.2.23, we design a polynomial f that is 0 at (x, y) when $x + y \in C$. The polynomial is the same as before, except that we include the factor $(x - y)$ to ensure that f is 0 when $x = y$, since $2x$ may not be in C. That is, let $f(x, y) = (x - y)\prod_{c \in C}(x + y - c)$. Note that $\deg(f) = m + 1$, where $m = |C|$.

We study the coefficient of $x^{a-1}y^{m-a+2}$. As before, contributions to the desired coefficient use x or y in each factor. The contributions choosing x in the first factor are positive, and those choosing $-y$ are negative. Thus $\left[x^{a-1}y^{m-a+2}\right]f(x, y) = \binom{m}{a-2} - \binom{m}{a-1} = \left[1 - \frac{a-1}{m-a+2}\right]\binom{m}{a-2}$. If $m \le 2a - 4$, then this coefficient is positive, and $a > m - a + 2$. Now the Combinatorial

Nullstellensatz guaratees $(x, y) \in A^2$ such that $f(x, y) \neq 0$. These are distinct elements of A whose sum is not in C, which is a contradiction. We conclude that $m \geq 2a - 3$. ■

The theorem below extends Theorem 20.2.24 to restricted sums over many variables. See Exercises 14–16 for the proof and applications.

20.2.25. THEOREM. (Alon–Nathanson–Rusza [1996]) Let p be a prime, and let h be a polynomial in k variables over $\mathbb{Z}_p$. Let $A_1, \dots, A_k$ be nonempty subsets of $\mathbb{Z}_p$, with $c_i = |A_i| - 1$ for all i. Let $m = \sum_{i=1}^{k} c_i - \deg(h)$. Let $C = \{\sum_{i=1}^{k} a_i \colon a_i \in A_i \text{ and } h(a) \neq 0\}$. If $\left[\prod_{i=1}^{k} x_i^{c_i}\right](\sum_{i=1}^{k} x_i)^m h(x) \neq 0$, then $|C| \geq m + 1$ (so $m < p$). ■

Our next consequence can also be considered number-theoretic, but it has a geometric application. It was conjectured by Artin [1934], proved by Chevalley [1936], and extended by Warning [1936]. The proof depends heavily on Fermat's Little Theorem (Corollary 0.@), which states that if p is a prime, then $a^{p-1} \equiv 1 \pmod{p}$ for every integer a not divisible by p.

20.2.26. THEOREM. (**Chevalley–Warning Theorem**) Let $P_1, \dots, P_m$ be polynomials over $\mathbb{F}_p$ in n variables. If $\sum_{i=1}^{m} \deg(P_i) < n$ and the polynomials share a zero, then they share another zero.

Proof: Let $(c_1, \dots, c_n)$ be a common zero. Let

$$f(x) = \prod_{i=1}^{m}(1 - P_i(x)^{p-1}) - \prod_{j=1}^{n}(1 - (x_j - c_j)^{p-1}).$$

Note that $f(c) = 1 - 1 = 0$. If there is no other common zero, then for $x \in \mathbb{F}_p^n - \{c\}$, there exists i such that $P_i(x) \not\equiv 0 \pmod{p}$, and there exists j such that $x_j \neq c_j$. By Fermat's Little Theorem, $P_i(x)^{p-1} \equiv 1 \equiv (x_j - c_j)^{p-1} \pmod{p}$. Hence $f(x) = 0$, for all $x \in \mathbb{F}_p^n$.

The degree of the first term in f is bounded by $(p-1)\sum_{i=1}^{m} \deg(P_i)$, which is less than $(p-1)n$. The degree of the second term is $(p-1)n$, and indeed $\left[\prod x_j^{p-1}\right] = (-1)^{n+1} \not\equiv 0 \pmod{p}$. Since $|\mathbb{F}_p| > p - 1$ and we choose each x_i from $\mathbb{F}_p$, Theorem 20.2.22 guarantees $x \in \mathbb{F}_p^n$ such that $f(x) \neq 0$. This contradiction proves that there must be another zero. ■

Chevalley proved this for $m = 1$ and Warning extended it; both guaranteed p solutions, which is stronger than proved here. We apply Theorem 20.2.26 to determine the transversal number of the hypergraph $\mathcal{H}$ of all hyperplanes in $\mathbb{F}_p^n$. The vertex set is $\mathbb{F}_p^n$, and for each hyperplane H we include an edge consisting of all the points in H. These points are the solutions to $a \cdot x = b$, for some $a \in \mathbb{F}_p^n$ and $b \in \mathbb{F}_p$. Every edge has p^{n-1} vertices. The **transversal number** $\tau(\mathcal{H})$ of a hypergraph $\mathcal{H}$ is the minimum size of a vertex subset intersecting all the edges.

20.2.27. THEOREM. (Jamison [1977], Brouwer–Schrijver [1978]) The transversal number of the hypergraph of all hyperplanes in $\mathbb{F}_p^n$ is exactly $n(p-1) + 1$.

Proof: First we produce a transversal of this size. Let B be the set of points in $\mathbb{F}_p^n$ having at most one nonzero coordinate; by construction B has the specified size. To prove that B is a transversal, we use induction on n. For $n = 1$, each point is a hyperplane, and indeed $B = \mathbb{F}_p^1$.

For $n > 1$, hyperplanes of the form $x_n = c$ are hit by the point in B having c in the last coordinate. For other hyperplanes, consider the fixed hyperplane H defined by $H = \{x \in \mathbb{F}_p^n \colon x_n = 0\}$. The hyperplanes of the form $x_n = c$ include H and all hyperplanes disjoint from H. The others intersect H in a hyperplane of $\mathbb{F}_p^{n-1}$ obtained by dropping the last coordinate (0) from the points in the intersection. By the induction hypothesis, these hyperplanes are hit by the points in B that have 0 in the last coordinate.

For the lower bound, let B be an arbitrary transversal. By applying a translation in each coordinate, we may assume that $0 \in B$. Let $A = B - \{0\}$. The set A intersects all hyperplanes not containing). This means that for all $x \in \mathbb{F}_p^n - \{0\}$, the equation $x \cdot y = 1$ has a solution $y \in A$.

Let $f(x) = \prod_{a \in A}(x \cdot a - 1)$. Since $x \cdot y = 1$ has a solution in A when $x \neq 0$, we have $f(x) = 0$ for $x \in \mathbb{F}_p^n - \{0\}$ and $f(0) = 1$. Now define a single polynomial P in $n(p-1)$ variables consisting of $p-1$ copies of x; that is $x^{(j)} = x_1^{(j)}, \dots, x_n^{(j)}$. Let $P = \left(\sum_{j=1}^{p-1} f(x_1^{(j)}, \dots, x_n^{(j)})\right) - (p-1)$.

Since f takes only the values 0 and 1, the sum is $p-1$ only when each summand is 1, so $x^{(j)} = 0$ for each j. Viewed over all $n(p-1)$ variables, 0 is the only zero. The contrapositive of the Chevalley–Warning Theorem (for $m = 1$) now yields $n(p-1) \leq \deg P \leq \deg f \leq |A| = |B| - 1$. ■

SUBGRAPHS WITH SPECIAL PROPERTIES

For our initial applications of the Combinatorial Nullstellensatz to graph theory, we present applications that we hope give further insight into how one goes about modeling a problem with a polynomial that will be permit application of the method.

Berge and Sauer conjectured that every 4-regular graph has a 3-regular subgraph. Taśkinov [1982] proved the conjecture. The claim is false for multigraphs (consider a 3-vertex graph with edges of multiplicity 2), but the conclusion becomes true when there is at least one "extra" edge, as seen by setting $p = 3$ in the next theorem. For convenience, when v is a vertex in a graph, we let $\Gamma(v)$ denote the set of edges incident to v.

20.2.28. THEOREM. (Alon–Friedland–Kalai [1984]) If p is prime, then every loopless multigraph G with average degree greater than $2p-2$ and maximum degree at most $2p-1$ contains a p-regular subgraph.

Proof: Suppose that G has n vertices and m edges. We want to design a function f such that when $f(x) \neq 0$, the point x selects for us a p-regular subgraph. Hence we introduce a variable x_e for each edge e, and we let $S_e = \{0,1\}$. To apply the Combinatorial Nullstellensatz, we will want a multilinear monomial term with a nonzero coefficient. Define f by

$$f(x) = \prod_{v \in V(G)} \Big[1 - \Big(\sum_{e \in \Gamma(v)} x_e\Big)^{p-1}\Big] - \prod_{e \in E(G)} (1 - x_e).$$

Each factor in the first term has degree $p-1$, so the degree of the first term is at most $(p-1)n$. This quantity is less than m, since the average degree exceeds $2p-2$. Hence the degree is determined by the second term, which has degree m, with $\big[\prod_{e \in E(G)} x_e\big] f(x) = (-1)^{m+1} \neq 0$.

By the Combinatorial Nullstellensatz, $f(\hat{x}) \neq 0$ for some $\hat{x} \in \{0,1\}^m$. Since $f(0) = 1 - 1 = 0$, this occurs with $\hat{x} \neq 0$. Since $\hat{x} \neq 0$, the second term in $f(\hat{x})$ has a factor that is 0. Hence the first term in $f(\hat{x})$ must be nonzero. By Fermat's Little Theorem, this requires that $\sum_{e \in \Gamma(v)}$ is a multiple of p for every vertex v.

Therefore, the degree of each vertex in the subgraph H of G with edge set $\{e \in E(G) \colon \hat{x}_e = 1\}$ is a multiple of p. Since $\Delta(G) \leq 2p-1$, the degree is always 0 or p. Since $\hat{x} \neq 0$, the degree is not always 0. Thus H has a nontrivial component, and it is a p-regular subgraph of G. ■

Theorem 20.2.28 leads to a bound on the number of edges needed in an n-vertex graph to ensure existence of a k-regular subgraph. Step 1 of the proof employs elementary observations in graph theory that we summarize briefly. The bound is not too far from optimal, as Pyber–Rödl–Szemerédi [1995] proved using probabilistic arguments that there are graphs with at least $\Omega(n \log\log n)$ that have no 3-regular subgraph. (They also showed that $O(n \log \Delta(G))$ edges force a 3-regular subgraph.)

20.2.29. THEOREM. (Pyber [1985]) If an n-vertex graph G has at least $32k^2 n \ln n$ edges, then G has a k-regular subgraph.

Proof: Let d be the average vertex degree in G. Say that an X, Y-bigraph is r-**halfregular** if $|X| \geq |Y|$ and every vertex of X has degree r.

Step 1. *Every graph with average degree d contains an r-halfregular X, Y-bigraph with $r \geq d/4$.* A bipartite subgraph with the most edges captures at least half the degree at each vertex, and the degree-sum formula then implies that it has at least half the edges and hence average degree at least $d/2$. Then, deleting a vertex v from a graph with average degree a increases the average degree if and only if v has degree less than $a/2$, so deleting vertices of minimum degree must eventually produce a subgraph with minimum degree at least $a/2$. Now we have a bipartite subgraph (an X, Y-bigraph with $|X| \geq |Y|$) with minimum degree r, where $r \geq d/4$. Obtain an r-halfregular X, Y-bigraph by deleting edges to reduce the degrees of vertices in X to r. Vertices in Y may wind up with smaller degree, but the average degree in Y will be at least r.

Step 2. *Every r-halfregular X, Y-bigraph H' contains an r-halfregular S, T-bigraph G' that has a perfect matching F'.* Choose S to be a minimal nonempty subset of X subject to $|N(S)| \leq |S|$. Such a set exists, since $|X| \geq |Y|$. If strict inequality holds, then deleting one element of S yields a smaller such set. Hence equality holds. By the minimality of S, always $|N(S')| > |S'|$ for $S' \subset S$. Hence Hall's Condition holds for the subgraph G' induced by $S \cup N(S)$, and G has a 1-factor. Since G' contains all edges incident to S, G' is r-halfregular.

Step 3. *An r-halfregular X, Y-bigraph H_0 contains edge-disjoint 1-regular subgraphs $F_0, \ldots, F_{r-1}$ with $V(F_0) \supseteq V(F_1) \supseteq \cdots \supseteq V(F_{r-1}) \neq \varnothing$.* Step 2 provides an r-halfregular subgraph G_0 of H_0 with 1-factor F_0. Having constructed $H_{i-1}, G_{i-1}, F_{i-1}$ with H_{i-1} being $(r-i+1)$-halfregular, let $H_i = G_{i-1} - E(F_{i-1})$. Since F_{i-1} is a 1-factor in G_{i-1}, the graph H_i is $(r-i)$-halfregular, and Step 2 applies to find G_i and F_i as desired.

Step 4. *Let p be a prime with $2k > p \geq k$. If $r \geq d/4 > 4p^2 \ln n$, then some $2p-1$ consecutive members of $F_0, \ldots, F_{r-1}$ together form a graph with average degree at least $2p-2$.* Let $n_j = |V(F_j)|$, and let $\hat{F}_j = \bigcup_{i=j}^{j+2p-2} F_i$. Since the number of edges in F_i is half its number of vertices, each of the matchings in $\hat{F}_j$ contributes at least $\frac{1}{2} n_{j+2p-2}$ edges. Being $2\left|E(\hat{F}_j)\right|/n_j$, the average degree is at least $2(2p-1)\frac{1}{2}\frac{n_{j+2p-2}}{n_j}$.

If $\frac{n_{j+2p-2}}{n_j} > \frac{2p-2}{2p-1}$ for some j, then the claim holds. Otherwise, $\left\lfloor \frac{r-1}{2p-2} \right\rfloor$ successive jumps of $2p-2$ steps yield

$$2 \leq n_{r-1} < n_0 \left(\frac{2p-2}{2p-1}\right)^{\frac{r-1}{2p-2}} \leq n \left(\frac{2p-2}{2p-1}\right)^{\frac{r-1}{2p-2}}.$$

Since $(2p-2)/(2p-1) = (1 - \frac{1}{2p-1}) < \mathrm{e}^{-1/(2p-1)}$, the upper bound simplifies to a negative power of n when $r > 4p^2 \ln n$, which contradicts that 2 is a lower bound.

Step 5. *If G has at least $32k^2 n \ln n$ edges, then G has a k-regular subgraph.* With this many edges, the average degree is at least $64k^2 \ln n$, and hence G has average degree at least $16p^2 \ln n$. By Step 1, G has an r-halfregular subgraph with $r \ge 4p^2 \ln n$. By Step 4, G has a subgraph with average degree greater than $2p-2$ and maximum degree at most $2p-1$. By Theorem 20.2.28, G has a p-regular subgraph. Since $p \ge k$ and we have arranged that this subgraph is bipartite, we can delete 1-factors from it to obtain a k-regular subgraph. ■

Instead of specifying the exact degree at each nonisolated vertex, we may be more flexible. Suppose that for each $v \in V(G)$ a *bad set* $B(v) \subseteq \{1,\dots,d_G(v)\}$ is specified. We seek a subgraph H such that $d_H(v) \notin B(v)$ for all v. Shirazi–Verstraëte [2008] gave an easy proof from the Combinatorial Nullstellensatz that there is a nontrivial such subgraph H when $\sum_{v\in V(G)} B(v) < |E(G)|$ (Exercise 23), and this is sharp.

They also proved a conjecture of Addario-Berry–Dalal–Reed–Thomason [2005] that allows 0 to be in the forbidden sets. This was stated originally in terms of "allowed" degrees, but it is a bit cleaner for bad degrees.

It is helpful to think in advance about the design of the polynomial f. We want the multivariate point x with $f(x) \ne 0$ to select the desired subgraph H. Hence we make a variable for each edge, and we allow it the values 0 and 1 to take model whether the edge is used in H. For each vertex v, we design a factor that is 0 when the constraint at v is violated.

20.2.30. THEOREM. (Shirazi–Verstraëte [2008]) For each vertex v in a graph G, specify a *bad set* $B(v) \subseteq \{0,\dots,d_G(v)\}$. If $|B(v)| \le \lfloor d(v)/2 \rfloor$ for all $v \in V(G)$, then G has a subgraph H with $d_H(v) \notin B(v)$ for all v.

Proof: Let $\Gamma(v)$ denote the set of edges in G incident to vertex v. Introduce a variable x_e for each edge e in G, and consider $x \in \{0,1\}^m$, where $m = |E(G)|$. Define a real-valued polynomial f by

$$f(x) = \prod_{v\in V(G)} \prod_{c\in B(v)} \Bigl(\sum_{e\in\Gamma(v)} x_e - c \Bigr).$$

Since $\sum_{e\in\Gamma(v)} x_e$ is the degree at v in a candidate subgraph, the factor for v is 0 if and only if the degree at v is forbidden. Hence we seek $x \in \{0,1\}^m$ such that $f(x) \ne 0$.

Since f is a product of linear factors, $\deg(f)$ is bounded by $\sum_{v\in V(G)} |B(v)|$. By the Combinatorial Nullstellensatz, it suffices to find a monomial with this degree having nonzero coefficient, whose variables all have exponent at most one. Monomials in the product arise by choosing, for each forbidden degree at each vertex, an edge incident to that vertex. We must not choose a given edge from both endpoints.

To avoid repeated selection, we orient G and pick for the monomial at vertex v only variables x_e such that v is the tail of e in the orientation. If the orientation has at least $\lfloor d(v)/2 \rfloor$ edges leaving each vertex v, then there are enough such edges to choose distinct ones for the elements of $B(v)$, since $|B(v)| \le \lfloor d(v)/2 \rfloor$. To form an orientation D such that $d_D^+(v) \ge \lfloor d_G(v)/2 \rfloor$, simply add a vertex w adjacent to all vertices of odd degree in G and orient by following an Eulerian circuit in each component.

We thus obtain a linear monomial. Every contribution to the coefficient of a monomial with degree $\sum_{v\in V(G)} |B(v)|$ is positive, since obtaining that degree requires selecting some x_e (and not c) from each factor.

Finally, when x is the point with $f(x) \ne 0$ guaranteed by the Combinatorial Nullstellensatz, each factor is nonzero, so the number of edges selected at v (via $x_e = 1$) does not lie in the bad set $B(v)$. ■

Theorem 20.2.30 is sharp, as the conclusion may fail as soon as one bad set is a bit too large. Let $G = K_{2r,2r}$, with partite sets X and Y. If $B(x) = \{0,\dots,r-1\}$ for $x \in X$ and $B(y) = \{r+1,\dots,2r\}$ for $y \in Y$, then each $B(v)$ has size $d(v)/2$, and a subgraph is good if and only if it is r-regular. As soon as the value r is added to one bad set, there is no longer a good subgraph.

Our next application comes from graph labeling and takes more work. The Combinatorial Nullstellensatz is a natural tool for graph labeling, because with variables for vertices the set S_i for variable x_i can list the labels allowed at vertex v_i.

20.2.31. REMARK. *Graph Labelings.* Ringel [1964] conjectured that K_{2m+1} decomposes into copies of any tree T with m edges. Attempts to prove this conjecture (or special cases) have tried to prove the stronger statement that there is a cyclically invariant decomposition, where the vertices are viewed as $\mathbb{Z}_{2m+1}$ and one copy of T is translated $2m+1$ times (note that K_{2m+1} has $(2m+1)m$ edges).

Rosa [1967] introduced several types of injective vertex labelings, called α-, β-, σ-, and ρ-"valuations". The conditions for a **ρ-valuation** are equivalent to cyclically invariant decomposition of K_{2m+1}; the vertices are assigned congruence classes modulo $2m+1$ so that the differences between adjacent labels are distinct elements of $[m]$. The conjecture that every tree with m edges has a ρ-valuation (and hence cyclically decomposes K_{2m+1}) is attributed to Kotzig [1973].

A **β-valuation** of a graph with m edges is an injective labeling of the vertices with integers in $\{0,\dots,m\}$ so that the differences of adjacent labels comprise $[m]$. Following Golomb [1972], a β-valuation is now known as a **graceful labeling**. A graceful labeling is a ρ-valuation in which the vertices are confined to $m+1$ consecutive congruence classes, so it yields a special type of cyclic decomposition. (The 5-cycle has a ρ-valuation but no graceful labeling.) The **Graceful Tree Conjecture**, also attributed to Kotzig, asserts that every tree has a graceful labeling.

An **α-valuation** is a special type of β-valuation with a value α such that the labels on each edge are on opposite sides of α; for bipartite graphs, this makes all the labels in one part bigger than all the labels in the other. The tree obtained by subdividing each edge of $K_{1,3}$ has no α-valuation. The "dynamic survey" by Gallian [2008] collects hundreds of results on these and other graph labelings. ■

Although most of the attention has been given to the Graceful Tree Conjecture, the conjecture that every tree has a ρ-valuation is presumably easier and still implies Ringel's decomposition conjecture. Kézdy [2006] used the Combinatorial Nullstellensatz to guarantee ρ-valuations for a special family of trees under the additional condition that the number $2m+1$ of congruence classes is prime.

20.2.32. DEFINITION. A tree with vertex set $\{v_0, \dots, v_m\}$ is **stunted** if it can be grown from the root vertex v_0 by successively introducing v_j for $1 \le j \le m$ so that v_j is made adjacent to some v_i with $i < j/2$.

In a stunted tree, the first two edges are incident to v_0, the third may be incident to v_0 or v_1, etc. The diameter of a stunted tree with m edges is at most $2\lg m$, whereas most trees have diameter around $\sqrt{m}$, but the number of leaves, the diameter, and the distance of leaves from the longest path may all be unbounded. In these senses the family is richer than others that are known to have ρ-valuations.

Given the canonical order in which the tree is grown, we index the edges as $e_1, \dots, e_m$ so that e_j joins v_j to an earlier vertex. With $\rho\colon V(T) \to \mathbb{Z}_{2m+1}$, we write the induced edge difference canonically as $\rho(v_j) - \rho(v_i)$, where $i < j$. The condition for ρ-valuation is then that these labels associated with edges are distinct and that none is the negative of another. Our polynomial f will incorporate a factor that models these requirements.

20.2.33. THEOREM. (Kézdy [2006]) If T is a stunted tree with m edges, and $2m+1$ is prime, then T has a ρ-valuation.

Proof: Let $p = 2m+1$. Let $P(j)$ be the set of indices of edges along the unique v_0, v_j-path in T; note that $P(0) = \varnothing$.

We associate variables with edges rather than vertices because translation by a constant does not change whether a labeling is a ρ-valuation,

and thus there are only m choices to make. Associating variable x_i with e_i, let $g_j(x) = \sum_{i \in P(j)} x_i$. With $P(0) = \varnothing$, actually $g_0(x) = 0$. Thus $g_j(x)$ can be viewed as $\rho(v_j)$, so that $g_j(x) - g_i(x) = \rho(v_j) - \rho(v_i)$. Define f by

$$f(x) = \prod_{1 \le i < j \le m} (x_j^2 - x_i^2) \prod_{0 \le i < j \le m} (g_j(x) - g_i(x)).$$

We have $f(x) \ne 0$ if and only if both factors are nonzero. With $\rho(v_j) = g_j(x)$ the second factor is nonzero if and only if the labeling is injective. Since by construction the values x_j are then the edge labels, the first factor is nonzero if and only if the edge labels are distinct and no two sum to 0. Hence $f(x)$ is nonzero for some $x \in \mathbb{Z}_p^m$ if and only if T has a ρ-valuation.

The set from which we choose each x_j has size $2m+1$, so the Combinatorial Nullstellensatz could apply to polynomials of degree up to $2m^2$. The degree of f is $2\binom{m}{2} + \binom{m+1}{2}$, so there is room to add helpful factors. Let $F(x) = f(x)\prod_{i=1}^m x_i^i$. The extra factor adds degree $\binom{m+1}{2}$, making the total degree $2m^2$. If $F(x) \ne 0$, then $f(x) \ne 0$, so it suffices to show that $F(x) \ne 0$ for some $x \in \mathbb{Z}_p^m$. By the Combinatorial Nullstellensatz, it suffices to show that $\left[\prod_{i=1}^m x_i^{2m}\right] F(x)$ is nonzero.

By Vandermonde's Identity,

$$\prod_{1 \le i < j \le m} (x_j^2 - x_i^2) = \sum_{\pi \in \mathbb{S}_m} \operatorname{sign}(\pi) \prod_{k=1}^m x_{\pi(k)}^{2(m-k)}.$$

Let $Q = \prod_{0 \le i < j \le m} (g_j(x) - g_i(x))$ and $R = \prod_{i=1}^m x_i^i$. To contribute to the coefficient of $\prod_{i=1}^m x_i^{2m}$, the term for π in the Vandermonde expansion must be multiplied by a term in QR with the factors $\prod_{k=1}^m x_{\pi(k)}^{2k}$. It thus suffices to show that the only nonzero term with exponents $2k$ for $1 \le k \le m$ in the expansion of QR is $\prod_{k=1}^m x_k^{2k}$.

Let $Q_{i,j} = g_j(x) - g_i(x)$. The factor $Q_{i,j}$ in Q is linear and equals $\sum_{k \in P(j)} x_k - \sum_{k \in P(i)} x_k$. To each term in the expansion of Q it contributes the factor $\pm x_k$ for some edge e_k on the v_i, v_j-path in T. The total degree for each term in Q equals $\sum_{i=1}^m i$.

Let $M = \prod_{k=1}^m x_{\sigma(k)} x^{2k}$. We claim that if M occurs in the expansion of QR, then M arises by selecting x_k from each linear factor $Q_{i,k}$ in Q, and hence $\sigma(k) = k$ for each k. Hence the desired monomial has coefficient ± 1.

We prove the claim by induction on k. The claim is vacuous for $k = 0$. For $k \ge 1$, note that $x_{\sigma(k)}$ has exponent $2k$ in M. For $j > 2k$, the exponent j on x_j in R is already too big. For $j < k$, the induction hypothesis yields $\sigma(j) = j$. Thus $k \le \sigma(k) \le 2k$.

Let $j = \sigma(k)$, and let T_i denote the subtree induced by $\{v_0, \dots, v_i\}$. Since $k \le j \le 2k$ and T is stunted, the graph obtained by appending e_j to T_{k-1} is also a tree (this is our only use of the hypothesis that T is stunted). For $0 \le i \le k-1$, the term $Q_{i,j}$ contributes the factor x_j to M, since by the induction hypothesis the terms for paths in T_i have selected the other variables as many times as they can be chosen. Thus x_j appears with exponent at least $k+j$ in M. We conclude that $\sigma(k) = k$, completing the proof. ■

It was not necessary to add the factor $\prod_{i=1}^{m} x_i^i$, and avoiding it would yield a ρ-valuation with more restricted choices for the labels, but using it simplified the argument.

THE ALON–TARSI THEOREM

In the extension of graph coloring known as list coloring, we still pick a single color for each vertex, but the set of colors available at each vertex may be restricted. When the colors represent resources, it may happen that some colors cannot be used at some vertices.

For example, when scheduling legislative committees during the week, two committees with a common member cannot meet at the same time. We form a graph with a vertex for each committee and an edge for each intersecting pair and seek a proper coloring. However, not all time slots are available for all committees, since the members of a committee may have prior commitments at some times. Each committee has a list of time slots when it can meet. How big do the lists need to be to guarantee finding time slots for all the committees without conflicts?

20.2.34. DEFINITION. A **list assignment** for a graph is a function L that assigns each vertex a list of available colors. The graph is **L-colorable** if it has a proper coloring f such that $f(v) \in L(v)$ for all v. Such a coloring is an **L-coloring**.

A graph G is **k-choosable** or **list k-colorable** if it is L-colorable whenever all lists have size at least k. The **list chromatic number**, **choice number**, or **choosability** $\chi_l(G)$ is $\min\{k\colon G \text{ is } k\text{-choosable}\}$.

The "lists" in list coloring are actually sets, without order. The term "list" is used by tradition and because "set coloring" refers to choosing sets of colors for the vertices. List coloring was introduced in Vizing [1976] and Erdős–Rubin–Taylor [1979]. Common notations for choice number include χ_l and ch. Edge-coloring also has a natural list version, with the edge-choosability $\mathrm{ch}'(G)$ being equal to $\mathrm{ch}(L(G))$.

The lists can be identical, so $\chi_l(G) \ge \chi(G)$. It is not possible to bound $\chi_l(G)$ in terms of $\chi(G)$; even for bipartite graphs it can be arbitrarily large (Exercise 31). On the other hand, many well-known upper bounds for $\chi(G)$ hold also for $\mathrm{ch}(G)$. It is easy to show that, $\mathrm{ch}(G) \le \Delta(G)+1$, and more generally every k-degenerate graph is $(k+1)$-choosable (`'Xdegen'`). One of the themes in 21st-century coloring theory is to strengthen upper bounds on $\chi(G)$ by showing that they are also upper bounds on $\mathrm{ch}(G)$ (see Exercises 29–).

For example, Brooks' Theorem extends to list coloring: When G is a connected graph, $\mathrm{ch}(G) \le \Delta(G)+1$ unless G is a complete graph or an odd cycle (Erdős–Rubin–Taylor [1979]). Galvin [1995] similarly strengthened the edge-coloring result for bipartite graphs: $\mathrm{ch}'(G) = \chi'(G)$ when G is bipartite. For planar graphs, one must give up a little: Thomassen [1994] proved the famous result that planar graphs are 5-choosable, and it was shown first by Voigt [1993] that this is sharp.

These and related results and extensions are studied at length in Section 3.4 of Volume I. Here our focus is to explore the natural use of the Combinatorial Nullstellensatz in problems involving list assignments.

Alon and Tarsi used a polynomial associated with a graph to obtain upper bounds on $\chi_l(G)$. We first state the result, which can be applied without knowing the algebraic background.

20.2.35. DEFINITION. A digraph D is a **circulation** if $d_D^+(v) = d_D^-(v)$ for all $v \in V(D)$. The parity of a circulation is the parity of the number of edges in it. Let $\mathrm{diff}(D)$ denote the absolute difference between the number of even circulations and the number of odd circulations contained in D (as spanning subgraphs).

In this context, what we call circulations have also been called "Eulerian digraphs" (see Alon [1993]). The term "circulation" reflects the lack of restriction on the number of components, and it generalizes for weighted graphs in the setting of network flows (see Chapter 7).

20.2.36. THEOREM. (Alon–Tarsi [1992]) Let $f(v) = 1 + d_D^+(v)$ for each vertex v in a digraph D. If $\mathrm{diff}(D) \ne 0$, then D is f-choosable. ■

20.2.37. Example. Let $G = C_n$. Let D be a cyclic orientation of G. Here $d_D^+(v) = 1$ for all v, and the only circulation contained in D are the trivial subgraph (no edges) and D itself. If n is even, then $\mathrm{diff}(D) = 2$, and G is 2-choosable. If n is odd, then $\mathrm{diff}(D) = 0$, and D gives us no information.

Now reverse one edge to form D'. The only circulation is the edgeless subgraph, and $\mathrm{diff}(D') = 1$. Since D' has a vertex with outdegree 2, we find that the odd cycle is 3-choosable. The theorem provides only upper bounds; we do not learn that the odd cycle is not 2-choosable.

An acyclic digraph contains only one circulation, with no edges. Thus Theorem 20.2.36 implies that G is k-choosable if G has an acyclic orientation in which every vertex has outdegree less than k. This again proves the trivial statement that k-degenerate graphs are $(k + 1)$-choosable. ■

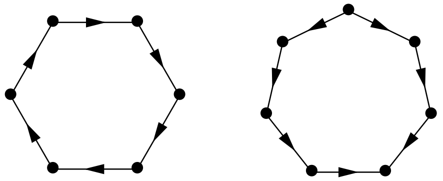

Before proving Theorem 20.2.36, we motivate it by describing several applications. A relatively easy application (Alon–Tarsi [1992]) is that every planar bipartite graph is 3-choosable (Exercise 88).

A most impressive application is the **Cycle-plus-triangles Theorem**. Consider a 4-regular graph formed from C_{3m} by adding m pairwise disjoint triangles. Du–Hsu–Hwang–Erdős [1987] conjectured that every such graph is 3-colorable. Fleishner–Stiebitz [1992] proved this by using Theorem 20.2.36 to prove the stronger result that every such graph is 3-choosable. Later, Sachs [1993] gave a combinatorial proof of 3-colorability.

The analysis of circulations in the Cycle-plus-triangles Theorem is lengthy. Instead we present an easier application to illustrate the Alon–Tarsi Theorem. Like the Cycle-plus-triangles Theorem, it proves 3-choosability for a 4-regular graph.

Recall that C_n^2 is the graph defined on n vertices around a circle by making each vertex adjacent to the four nearest vertices; it is 4-regular. If we can orient it with two edges in and two out at each vertex such that the numbers of circulations of even and odd size differ, then it is 3-choosable. The motivation was in the context of "total coloring".

20.2.38. Example. Total coloring is coloring the vertices and the edges so that no adjacent or incident objects have the same color (see Exercises 52–53). The **total chromatic number** is the minimum number of colors needed; it is the chromatic number of the **total graph**, obtained from a graph G by subdividing every edge and then taking the square of the graph (adding edges joining vertices at distance 2 in the subdivision graph). For a cycle, subdividing just doubles the length, so total coloring of a cycle is proper coloring of the square of a cycle twice as long.

Note that C_n^2 is not 3-colorable when n is not divisible by 3, so $3 \mid n$ is certainly a necessary condition for 3-choosability. ■

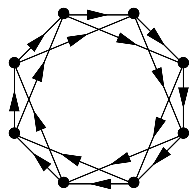

20.2.39. THEOREM. (Juvan–Mohar–Škrekovski [1998]) The graph C_n^2 is 3-choosable if and only if n is divisible by 3. Consequently, C_m is 3-total-choosable when m is divisible by 3.

Proof: We observed that $3 \mid n$ is necessary. For the converse, we seek a suitable orientation of C_n^2 with maximum outdegree 2. Form D by orienting every edge in the clockwise direction as shown above.

Among the circulations in D, let S_0, S_1, S_2 be the sets of circulations in which all vertices have outdegree at most 1, all equal 1, or all at least 1, respectively. Obviously, $S_0 \cap S_2 = S_1$.

Less obviously, $S_0 \cup S_2$ is the set of all circulations. Let D' be a circulation in D. If $d^+_{D'}(v) = 0$, then at most one edge of D' "crosses the gap" from v to the next vertex along the circle. Whenever D' has at most one edge crossing such a gap, also D' has at most one edge crossing the next gap. Hence we cannot encounter a vertex with outdegree 2 in D'.

To prove $\mathrm{diff}(D) \neq 0$, we prove that the total number of circulations is congruent to 2 modulo 4 and that the numbers of even circulations and odd circulations are both even. The latter statement is easy. For a circulation H, the remaining edges in D also form a circulation H'. Although $H \neq H'$, the numbers of edges in H and H' have the same parity, since D has $2n$ edges. Hence each parity class of circulations has even size.

Now consider all circulations. In S_1, there are exactly two: the "outside" edges of length 1 and the "inside" edges of length 2; the latter may be one cycle or two depending on the parity of n. Furthermore, complementation of edge sets matches $S_0 - S_1$ with $S_2 - S_1$; they have the same size. Hence it suffices to show that $|S_0 - S_1|$ is even.

We claim that $|S_0 - S_1|$ is the number of cycles that wrap once around the circle, with steps whose lengths total n. One cycle (n unit steps) is an element of S_1, while one element of $S_0 - S_1$ (the circulation with no edges) is not counted among the cycles. When n is odd, the n steps of length 2 form a cycle, but it is neither in $S_0 - S_1$ nor in the set of cycles we count.

For a fixed vertex v, a cycle that omits v corresponds to a $1,2$-list with sum $n-2$, while a cycle that vists v corresponds to a $1,2$-list with sum n. The adjusted Fibonacci number $\hat{F}_i$ is the number of $1,2$-lists with sum i. We obtain $|S_0 - S_1| = \hat{F}_{n-2} + \hat{F}_n = \hat{F}_{n-1} + 2\hat{F}_{n-2}$. Note that $\hat{F}_0 = \hat{F}_1 = 1$, and hence $\hat{F}_2$ is even. The parity pattern then repeats, with $\hat{F}_r$ even if and only if $r \equiv 2 \pmod 3$). Thus $|S_0 - S_1|$ is even if and only if $3 \mid n$. ■

Juvan–Mohan-Škrekovski [1998] stated the result only for $6 \mid n$, since that is when it corresponds to total coloring of cycles. The key idea is used also in the proof of the Cycle-plus-triangles Theorem and in many applications of the Alon–Tarsi Theorem to 4-regular graphs: show that in the specified orientation the total number of circulations is an odd multiple of 2 and the number with even size is even.

Theorem 20.2.36 has also been applied to edge-choosability of planar graphs. Jaeger and Tarsi independently observed (see Alon [1993]) that every 2-connected 3-regular planar graph is 3-edge-choosable (using also the Four Color Theorem). Ellingham and Goddyn [1996] extended this, proving that every k-regular k-edge-colorable planar graph is k-edge-choosable. Thus the List Coloring Conjecture holds for these graphs.

Now we develop the proof of Theorem 20.2.36. We use an algebraic interpretation of L-coloring when the colors are real numbers. Associate with each vertex v_i a variable x_i. We define a polynomial that is nonzero just when the numbers assigned to the vertices form a proper coloring.

20.2.40. DEFINITION. Given a graph G with vertex set $v_1, \ldots, v_n$, let $E'(G) = \{(i, j)\colon i < j \text{ and } v_iv_j \in E(G)\}$. The **graph polynomial** p_G of G is defined by $p_G(x_1, \ldots, x_n) = \prod_{(i,j)\in E'(G)}(x_i - x_j)$.

Early applications of the graph polynomial include Petersen [1891], Scheim [1974], Li–Li [1981]. Indeed, Petersen introduced graphs to study such polynomials (see Toft [1992]). Throughout this discussion, we assume that G is a graph with a fixed vertex indexing $v_1, \ldots, v_n$.

As remarked, $p_G(x_1, \ldots, x_n)$ is 0 if and only if assigning x_i to v_i for all i produces a monochromatic edge. The relation of p_G to orientations is seen by expanding the product. For a transitive tournament, parity of orientations reduces to parity of permutations.

20.2.41. DEFINITION. An edge v_iv_j in an orientation of G (with vertices $v_1, \ldots, v_n$) is *decreasing* if $i > j$. The **parity** of an orientation of G is the parity of the number of decreasing edges.

20.2.42. LEMMA. Let d denote a vector $d_1, \ldots, d_n$ of nonnegative integers. In the graph polynomial p_G, the coefficient of the monomial $\prod x_i^{d_i}$ is the number of even orientations of G with outdegrees d minus the number of odd orientations of G with outdegrees d.

Proof: The polynomial is homogeneous of degree $|E(G)|$, since each factor is homogeneous of degree 1. Each contribution to the expansion is formed by selecting one endpoint of each edge. This corresponds to an orientation by letting the selected vertex be the source of the edge. The resulting contribution to the expansion is $(-1)^t \prod x_i^{d_i}$, where d_i is the outdegree of v_i in the corresponding orientation and t is the number of decreasing edges. For a given list d of outdegrees, the even orientations count $+1$, and the odd orientations count -1. ■

In order to relate this coefficient to diff(D) in the statement of Theorem 20.2.36, we establish a bijection from the set of orientations with the same outdegrees as D to the set of circulations contained in D.

20.2.43. LEMMA. For an orientation D of G with $d_i = d_D^+(v_i)$ for each i, the absolute value of the coefficient of $\prod x_i^{d_i}$ in p_G is diff(D).

Proof: Fixing the orientation D, let D' be any orientation of G having outdegree at each vertex the same as in D. Let $D \oplus D'$ be the spanning subdigraph of D whose edges are those reversed in D'. Since $d_D^+(v) = d_{D'}^+(v)$ for all v, the subdigraph $D \oplus D'$ is a circulation contained in D. Also, each edge placed in $D \oplus D'$ changes the parity of the number of decreasing edges, so $D \oplus D'$ has an even number of edges if and only if the numbers of decreasing edges in D and D' have the same parity.

Switching the orientation on the edges in a circulation does not change the outdegree at any vertex, so this inverts the map sending D' to $D \oplus D'$. We thus have a bijection from the set of orientations with the same outdegrees as D to the set of circulations contained in D. We have also observed that the parities of an orientation and the corresponding circulation are the same. Hence the coefficient has the value claimed. ■

By Lemma 20.2.43, diff(D) depends only on the outdegrees in D. The Alon–Tarsi Theorem now follows immediately.

Proof of Theorem 20.2.36 (Alon–Tarsi Theorem). Given a graph G with vertices $v_1, \ldots, v_n$ and an orientation D of G with diff$(D) \neq 0$, let $d_i = d_D^+(v_i)$. By Lemma 20.2.43, $\left|\left[\prod x_i^{d_i}\right] p_G\right| = \text{diff}(D)$. Since diff$(D) \neq 0$, Theorem 20.2.22 implies that p_G is nonzero for some $x \in \prod S_i$ when $|S_i| \geq d_i + 1$ for each i. With $f(v) = 1 + d_D^+(v)$ for each v, we conclude that D (and the underlying graph G) is f-choosable. ■

Ramamurthi–West [2005] generalized the Alon–Tarsi Theorem to k-uniform hypergraphs, where k is prime. The notion of orientation is selecting a source from each edge. Like the graph polynomial, the hypergraph polynomial has a factor for each edge; it is $\sum_{i=0}^{k} \theta^i x_{j_i}$, where θ is a kth root of unity and the vertices $v_{j_1}, \dots, v_{j_k}$ are in increasing order of indices. The notion of circulation is more complicated, but the proof is a straightforward generalization of the Alon–Tarsi proof.

LIST WEIGHTING

The Combinatorial Nullstellensatz is well-suited for problems involving lists, where the sets S_i of values for the variables can be interpreted as lists from which choices are to be made. This allows us to generalize various labeling problems to a list context.

20.2.44. DEFINITION. A **weighting** f of a graph G assigns an integer weight to each edge of G and generates a color $\phi(v)$ at each vertex v that is the sum of the weights on edges incident to v. A **total weighting** f assigns weights to both vertices and edges, and then $\phi(v)$ is $f(v)$ plus the sum of the weights on the incident edges. In either case, f is **proper** if ϕ is a proper coloring of G.

We use “proper” since the weighting produces a proper vertex coloring; “neighbor-distinguishing” was an earlier term for this. The motivation for the problem is that when the weights are taken as edge-multiplicities, adjacent vertices will have distinct degrees. It is then natural to seek such a weighting with small multiplicities. Note that K_2 has no proper weighting, although it has a proper total weighting.

20.2.45. CONJECTURE. (**1,2,3-Conjecture**; Karonski–Łuczak–Thomason [2004]) Every graph without isolated edges has a proper weighting using weights in $\{1, 2, 3\}$.

20.2.46. Example. *K_n has no proper $1, 2$-weighting.* In a regular graph G, subtracting a constant from all edge weights does not affect distinctness of sums at vertices. Hence we may seek a proper weighting from $\{0, 1\}$, which is equivalent to a subgraph H (the edges of weight 1) such that adjacent vertices of G have distinct degrees in H. This fails when $G = K_n$, since every n-vertex graph has two vertices of the same degree, by the pigeonhole principle (degrees 0 and $n - 1$ cannot both occur). ■

The first step was to show that bounded multiplicity always suffices, equivalent to proper weighting from the set $[k]$. Addario-Berry, Dalal, McDiarmid, Reed, and Thomason [2007] showed that $k = 30$ suffices. This was reduced to $k = 16$ in Addario-Berry–Dalal–Reed [2008] and to $k = 13$ in Karonski–Łuczak–Thomason [2004]. A subsequent breakthrough reduced the bound to $k = 5$ (Kalkowski–Karoński–Pfender [2010]). The proof of this result was motivated by a breakthrough on the analogous conjecture about total weightings.

20.2.47. CONJECTURE. (**1,2-Conjecture**; Przybyło–Woźniak [2010]) Every graph has a proper total weighting using weights in $\{1, 2\}$.

Przybyło and Woźniak verified the conjecture for complete graphs (Exercise 20), 4-regular graphs, and graphs with chromatic number at most 3. They also showed that weights in [11] always suffice. The breakthrough is the following theorem, with a remarkably simple proof.

20.2.48. THEOREM. (Kalkowski [2009+]) Every graph has a proper total weighting with vertex weights in $\{1, 2\}$ and edge weights in $\{1, 2, 3\}$.

Proof: Begin with two chips on each edge and one chip on each vertex. The weight of a vertex or edge will be the final number of chips on it.

Process the vertices in some order $v_1, \dots, v_n$. We move a chip onto or off an edge only when we process its later endpoint. When processing v_i, we make its color (the number of chips on it and its incident edges) different from the colors of its earlier neighbors, and its color never changes thereafter. Doing this for $v_1, \dots, v_n$ completes the proof.

Before processing v_i, the edges from v_i to its d earlier neighbors have two chips. If an earlier neighbor x has one chip, we may move one chip from $v_i x$ to x. If x has two chips, then we may move one chip from x to $v_i x$. In each case, the number of chips on x remains in $\{1, 2\}$, the number of chips on xv_i remains in $\{1, 2, 3\}$, and *the color of x does not change*.

For each back edge $v_i x$, the two possibilities make contributions to the color at v_i that differ by 1 (also there is 1 for v_i and 2 for each edge to later neighbors). The difference between taking each lower option and taking each higher option is d, and every value between the lowest and the highest is achievable. Altogether, $d+1$ consecutive values are achievable for the color of v_i. Since the earlier neighbors occupy only d colors, we can shift chips between the back neighbors and back edges to give v_i a color different from its earlier neighbors. ■

The proof that graphs without isolated edges have proper weightings from $\{1, 2, 3, 4, 5\}$ has a similar flavor. Edges initially have three chips, and processing of back edges may be add or remove two chips. Since vertices have no chips, this changes the colors on earlier neighbors. Hence a stronger hypothesis is needed about the weights used on the earlier neighbors, and the option of removing or adding one chip is also needed.

List versions of these conjectures allow arbitrary lists of integer labels at the vertices or edges, from which the weights must be chosen.

20.2.49. DEFINITION. A (k, k')-**total list assignment** for a graph G is a map L that assigns each vertex a set of k numbers and each edge a set of k' numbers. The graph G is (k, k')-**weight-choosable** if for every (k, k')-total list assignment, a proper total weighting can be chosen from the lists.

Since a $(1, 3)$-total list assignment may assign list $\{0\}$ to each vertex, $(1, 3)$-weight-choosability is stronger than choosability when just edge weights from lists of size 3, which in turn is stronger than the 1,2,3-Conjecture. Bartnicki–Grytczuk–Niwczyk [2009] conjectured the intermediate property for every graph without isolated edges; they proved this for complete graphs, complete bipartite graphs, and others. In fact, their argument showed that these graphs are $(1, 3)$-weight-choosable. This suggests the strongest possible conjectures for weight choosability.

20.2.50. CONJECTURE. (Wong–Zhu [2009+]) Every graph is $(2, 2)$-weight-choosable. Every graph without isolated edges is $(1, 3)$-weight-choosable.

We now drop the word “weight” and just write (k, k')-choosable. To motivate this, note that a $(k, 1)$-total list assignment may assign list $\{0\}$ to every edge, and hence every $(k, 1)$-weight-choosable graph is k-choosable. Thus (k, k')-choosability generalizes the notion of k-choosability. Since larger lists don’t hurt, every (k, k')-choosable graph is $(k, k'+1)$-choosable and $(k+1, k')$-choosable. Thus a weaker conjecture is that for some (k, k') every graph is (k, k')-choosable. The conjectures are sharp, since some graphs are not $(1, 2)$-choosable.

20.2.51. Example. *Failure of* $(1, 2)$*-choosability.* For the complete graph K_n, we showed in Example 20.2.46 that no proper total-weighting can be chosen when the vertex lists are all $\{0\}$ and the edge lists are all $\{0, 1\}$.

Our second example is the small tree with vertex lists of size 1 shown below and edge lists all $\{0, 1\}$. Since the central edge contributes equally to its endpoints, the choices on the two “ends” must contribute different amounts to the central vertices. If the pendant edges on one side contribute 0 and 1, then we will not be able to assign 0 or 1 to the middle edge. Hence we must have two 0s on one side and two 1s on the other, but now selecting 0 for the middle 0 violates the edges on the side with 1s, and selecting 1 violates the edges on the side with 0s. ■

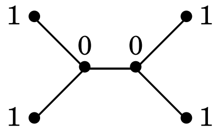

20.2.52. REMARK. The tree in Example 20.2.51 has an odd number of edges. Infinitely many trees of odd size fail to be $(1, 2)$-choosable (Exercise 21), but Wong–Zhu [2009+] used the Combinatorial Nullstellensatz to show that all trees with an even number of edges are $(1, 2)$-choosable. From this it follows easily that all trees are $(2, 2)$-choosable (Exercise 22).

Wong–Zhu [2009+] also applied the Nullstellensatz to show that complete graphs (and in general all complements of linear forests) are $(2, 2)$-choosable (true also for cycles, generalized theta graphs, etc.). Via an inductive proof (not Nullstellensatz), Wong–Yang–Zhu [2009+] proved $(1, 2)$-choosability for complete multipartite graphs with at most two parts of size greater than 2 (this includes complete bipartite graphs).

Later, Wong–Zhu [2012+] proved that every graph is $(2, 3)$-choosable. This statement does not prove either the $1, 2, 3$-Conjecture or the $1, 2$-Conjecture and is weaker than Conjecture 20.2.50, but it is stronger than Theorem 20.2.48 and applies to all graphs. The setup using the Combinatorial Nullstellensatz is the same as for their work on $(1, 2)$-choosability, but the combinatorial tricks to produce the nonzero coefficient are different. We will present both results. ■

20.2.53. REMARK. (k, k')*-choosability via the Combinatorial Nullestellensatz.* Let G be a graph with n vertices and m edges. To apply the Nullstellensatz, we need a polynomial f where the variables take values from the lists, and the value of f is nonzero if and only if the resulting weighting is proper. The choice of f is obvious. Given variables $x_1, \dots, x_m$ corresponding to the edges $e_1, \dots, e_m$ and variables $y_1, \dots, y_n$ corresponding to the vertices $v_1, \dots, v_n$, let

$$f(x, y) = \prod_{uw \in E(D)} (\phi(w) - \phi(u)),$$

where $\phi(v_i)$ is the total weight seen by vertex v_i and D is a fixed orientation of G, chosen just to specify the signs. The value of f is nonzero if and only if the weights given by the values of x and y form a proper total weighting of G, regardless of the choice of D. With $\Gamma(v)$ denoting the set of edges incident to a vertex v, we have $\phi(v_i) = y_i + \sum_{e_j \in \Gamma(v_i)} x_j$.

Each factor is a homogeneous polynomial of degree 1, so f has degree m. To prove (k, k')-choosability via the Combinatorial Nullstellensatz, we want the coefficient of an appropriate monomial in the expansion of f to be nonzero. This suffices if the degree of each variable in the monomial is less than the size of the corresponding list. For $(1, 2)$-choosability, there is only one choice: the term must be $\prod_{i=1}^{m} x_i$, with each edge-variable having degree 1 and each vertex-variable having degree 0. For $(2, 2)$-choosability, it suffices to have a nonzero term using any m variables from among the edges and vertices to have degree 1. For $(2, 3)$-choosability, we can allow edge variables to appear with degree 2. ∎

The next lemma describe the coefficients of the multilinear terms.

20.2.54. LEMMA. Let G be a graph with vertices $v_1, \ldots, v_n$ and edges $e_1, \ldots, e_m$. Given an orientation D, let $f(x, y)$ be the polynomial defined in Remark 20.2.53. Defina a matrix A with m rows indexed by edges and $m + n$ columns indexed by edges and vertices, by

$$A_{uw,z} = \begin{cases} -1 & \text{if } z = u \text{ or } z \in \Gamma(u) - \{uw\}, \\ 1 & \text{if } z = w \text{ or } z \in \Gamma(w) - \{uw\}, \\ 0 & \text{otherwise.} \end{cases}$$

If S is a multiset of m variables from $\{x_1, \ldots, x_m\} \cup \{y_1, \ldots, y_n\}$, with each variable z selected $h(z)$ times, then the coefficient of $\prod_{z \in S} z$ in $f(x, y)$ is $1/\prod_z h(z)!$ times the permanent of the matrix B whose columns are the columns of A for S (with multiplicity).

Proof: The factor in f for the edge uw in D is the dot product of the row in A indexed by uw with the vector $(x_1, \ldots, x_m, y_1, \ldots, y_n)$ of variables. A nonzero contribution to a coefficient in the expansion of f is the product of a term from each of the m factors. The selected term from a factor e is a variable for a column in A, and the sign on it is ± 1 as recorded in the row of A corresponding to e. Selecting a variable h times can be modeled by repeating the column h times among the m columns in B. However, each contribution that selects this variable h times appears $h!$ times in the computation of the permanent of B. The computation of $\operatorname{per} B$ sums all such contributions. ∎

In particular, we obtain a nonzero coefficient for the term in $f(x, y)$ with specified exponents if and only if the m-by-m matrix using columns of A with the specified multiplicities has nonzero permanent.

20.2.55. REMARK. Lemma 20.2.54 implies that the $(1, 3)$-choosability part of Conjecture 20.2.50 would follow from the **Permanent Conjecture** of J. Kahn: If A is an invertible m-by-m matrix, then there is an m-by-m submatrix of $[AA]$ having nonzero permanent. Yang Yu generalize the Permanent Conjecture to submatrices of $[AB]$ whenever A and B are invertible over the same field. The general conjecture is equivalent to the Additive Basis Conjecture of Alon and Tarsi and has various consequences in graph theory, including that every 6-edge-connected graph has a nowhere-zero 3-flow (see Chapter 9). ∎

20.2.56. THEOREM. (Zhu–Wong [2013+]) Every graph is $(2, 3)$-choosable.

Proof: We use induction on the number of vertices to construct the desired matrix with nonzero permanent from columns of the matrix A^G defined in Lemma 20.2.54. Let $A(z)$ denote the column of A^G corresponding to the variable for z. We use repeatedly the important observation that $A(uv) = A(u) + A(v)$ for each edge uv. Also important is the linearity of the permanent in each column: if matrices B and B' are the same except in column j, and C agrees with both except for its jth column being the sum of the other two, then $\operatorname{per} C = \operatorname{per} B + \operatorname{per} B'$.

The claim holds vacuously for K_1; let G be a larger graph. Let v be a vertex of G, and let $G' = G - v$. Let M' be the matrix for G' guaranteed by the induction hypothesis, using columns of $A^{G'}$ with appropriate multiplicity. Let $e_1, \ldots, e_d$ be the edges incident to v, oriented toward v. Index the other endpoints of $e_1, \ldots, e_d$ as $v_1, \ldots, v_d$, respectively. Let M be the matrix consisting of M' in the upper left, plus rows for $e_1, \ldots, e_d$ and d copies of the column for v in A^G; this column $A(v)$ is 0 in the rows corresponding to edges of G' and 1 in the rows for $e_1, \ldots, e_d$. Nonzero contributions to $\operatorname{per} M$ must take m factors from d-by-d matrix in the lower-right corner. Hence $\operatorname{per} M = d! \operatorname{per} M' \neq 0$.

$$M = \begin{array}{c|cccccc|}
 & & & & v & \cdots & v \\
 & & & & 0 & \cdots & 0 \\
 & & M' & & \vdots & & \vdots \\
 & & & & 0 & \cdots & 0 \\
e_1 & \cdot & \cdot & \cdot & 1 & \cdots & 1 \\
\vdots & \cdot & \cdot & \cdot & \vdots & & \vdots \\
e_d & \cdot & \cdot & \cdot & 1 & \cdots & 1
\end{array}$$

We now have a matrix built from columns of A^G with nonzero permanent, but it uses d copies of $A(v)$. Note first that since $A(v_i) = A(e_i) - A(v)$ for $1 \le i \le d$, we can think of any use of column $A(v_i)$ as yielding a sum of two permanents. The matrix with $A(v_i)$ replaced by $A(v)$ has $d + 1$

columns in which all the nonzero entries are confined to d rows; hence its permanent is 0. This means that in M we can replace the appearance of any column for a vertex v_i adjacent to v with the column for the corresponding edge e_i, without changing the permanent.

We are now free to reintroduce $A(v_i)$, and we are also free to use another copy of $A(e_i)$. Now treat the d copies of $A(v)$ as copies of $A(e_1) - A(v_1), \ldots, A(e_d) - A(v_d)$. Expanding via linearity expresses $\operatorname{per} M$ as the sum of 2^d permanents. Since $\operatorname{per} M \neq 0$, at least one of the resulting matrices has nonzero permanents and uses each column of A^G with sufficiently small multiplicity. ■

20.2.57. REMARK. If we replace all copies of $A(v)$ in the final step of the proof of Theorem 20.2.56, then $A(v)$ appears with multiplicity 0. If for some i we omit the substitution of $A(e_i) - A(v)$ for $A(v_i)$ in the first step and then also the substitution of $A(e_i) - A(v_i)$ for $A(v)$ in the last step, then $A(e_i)$ appears with multiplicity 0.

These observations imply that a proper total weighting can still be chosen from a total list assignment where the vertices of an independent set are given lists of size 1, or where the edges of a spanning tree are given lists of size 1. More generally, we can specify arbitrary weights at any independent set and on the edges of any forest not using those vertices, and still we can complete a proper total weighting from lists of size 2 on the remaining vertices and size 3 on the remaining edges. ■

Now we return to $1,2$-choosability for trees. As noted in Remark 20.2.53, for a proof of $(1,2)$-choosability we need to show that $\left[\prod_{i=1}^{m} x_i\right] f(x,y)$ is nonzero, where $x_1, \ldots, x_m$ are the variables for the edges. Gerard Chang noted a useful fact about this.

20.2.58. LEMMA. Let A^G be the matrix of Lemma 20.2.54 for the connected graph G with orientation D, and let $B(G)$ denote the square submatrix of A^G whose columns correspond to the edge variables. If G has a cut-edge e, with $G - e$ having components G_1 and G_2, then

$$\operatorname{per} B(G) = \operatorname{per} B(G_1) \operatorname{per} B(G_2') + \operatorname{per} B(G_1') \operatorname{per} B(G_2),$$

where G_i' is the subgraph consisting of G_i plus e (and its endpoint).

Proof: We apply row-linearity of the permanent to $B(G)$. Let $E(G_1) = \{e_1, \ldots, e_{t-1}\}$ and $E(G_2) = \{e_{t+1}, \ldots, e_m\}$, so $e = e_t$. Letting z denote an entry that may be anything in $\{0, 1, -1\}$, we have

$$\begin{array}{cc|ccccccc|}
 & & x_1 & \cdots & x_{t-1} & x_t & x_{t+1} & \cdots & x_m \\
 & e_1 & & & & z & & & \\
 & \vdots & & B(G_1) & & z & & 0 & \\
 & e_{t-1} & & & & z & & & \\
B(G) = & e_t & z & z & z & 0 & z & z & z \\
 & e_{t+1} & & & & z & & & \\
 & \vdots & & 0 & & z & & B(G_2) & \\
 & e_m & & & & z & & &
\end{array}.$$

Now obtain u^1 and u^2 from row t of $B(G)$ by setting the first $t-1$ positions or last $m-t$ positions to 0, respectively. Let B_i be the matrix obtained from $B(G)$ by replacing row t with u^i. Since row t of $B(G)$ is $u^1 + u^2$, we have $\operatorname{per} B(G) = \operatorname{per} B_1 + \operatorname{per} B_2$. To obtain a nonzero contribution to $\operatorname{per} B_1$, the nonzero terms from rows $e_t, \ldots, e_m$ lie in columns $x_t, \ldots, x_m$, so the nonzero terms from rows $e_1, \ldots, e_{t-1}$ must lie in columns $x_1, \ldots, x_{t-1}$. Thus $\operatorname{per} B_1 = \operatorname{per} B(G_1) \operatorname{per} B(G_2'$; similarly, $\operatorname{per} B_2 = \operatorname{per} B(G_1') \operatorname{per} B(G_2)$. ■

20.2.59. THEOREM. (Wong–Zhu [2009+]) Every tree with an even number of edges is $(1,2)$-choosable.

Proof: Let G be a tree with m edges. For $(1,2)$-choosability, by Remark 20.2.53 and Lemma 20.2.54 it suffices to prove that $\operatorname{per} B(G) \neq 0$. To prove this when m is even, we prove that $\operatorname{per} B(G) \equiv m-1 \pmod 2$. Since all nonzero contributions are ± 1 and we care only about parity, we may change all nonzero entries to $+1$ and ignore the orientation. The matrix $B(G)$, expressed modulo 2 in this way, becomes actually the adjacency matrix of the line graph of G.

We apply Lemma 20.2.58 to prove the claim inductively. For all of G_1, G_2, G_1', G_2' to be smaller than G, the cut-edge e in the induction step must not be incident to a leaf. Hence the base step is when G is a star.

When G is a star, each edge is incident to every other. Hence $B(G) = J_m - I_m$, where J_m and I_m are the m-by-m all-1 and identity matrices. The permanent equals the number of permutations of $[m]$ with no fixed point, which is the derangement number D_m. The derangement numbers are computed recursively by $D_0 = 1$, $D_1 = 0$, and $D_m = (m-1)(D_{m-1} + D_{m-2})$ (see Chapter 16). Reducing modulo 2, it follows immediately by induction that $D_m \equiv (m-1) \pmod 2$.

For the induction step, when G is not a star it has an edge e_t incident to no leaf, with $G - e_t$ having components G_1 and G_2 with $t-1$ and $m-t$ edges, respectively. Since each matrix in the formula of Lemma 20.2.58 is the special matrix that arises from a tree, Lemma 20.2.58 and the induction hypothesis yield

$$\begin{aligned}\operatorname{per} B(G) &= \operatorname{per} B(G_1)\operatorname{per} B(G_2') + \operatorname{per} B(G_1')\operatorname{per} B(G_2)\\ &\equiv (t-2)(m-t) + (t-1)(m-t-1)\\ &\equiv 2(t-2)(m-t-1) + (t-2) + (m-t-1) \equiv m-1 \pmod 2. \quad \blacksquare\end{aligned}$$

20.2.60. REMARK. In proving that K_n (and complements of linear forests) are $(2,2)$-choosable, again we use the same polynomial, producing the same matrix of coefficients. However, since vertex lists and edge lists both have size 2, we may consider any m variables, where $m = |E(G)|$. Choosing these columns yields a square matrix, and again the task is to show that the permanent is nonzero (not modulo 2). The difficulty is picking the right columns (variables) to delete.

Index the vertices as $v_1, \dots, v_n$. It turns out that the computation becomes feasible when the variables for vertex v_1 and edges $v_i v_{i+1}$ with $1 \le i \le n-1$ are deleted. The permanent of the matrix with these columns deleted is the coefficient for the multilinear monomial with the remaining m variables. Using induction and linearity properties of the permanent, one can show that this permanent is nonzero. ■

EXERCISES

20.2.1. (−) In a town with n people, there are m sports clubs $A_1, \dots, A_m$ and m theater clubs $B_1, \dots, B_m$. Prove that if $|A_i \cap B_i|$ is odd for every i, and $|A_i \cap B_j|$ is even whenever $i < j$, then $m \le n$, and this is sharp.

20.2.2. Prove that there are between $2^{n(n+2)/8}/(n!)^2$ and $2^{n^2}/n!$ nonisomorphic sets of n odd-sized subsets of $[n]$ such that the intersection of every pair has even size. Two such sets are isomorphic if one can be obtained from the other by permuting $[n]$. (Hint: Let $n = 2k$. From a k-by-k binary matrix A, form an n-by-n binary matrix $\left(\begin{smallmatrix} A+I_k & A\\ A & A+I_k\end{smallmatrix}\right)$.) (Szegedy [1988])

20.2.3. (◇) Let A be an m-by-m matrix of integers. Prove that if some prime power divides every off-diagonal entry but no diagonal entry, then A is nonsingular. Conclude that if the greatest common divisor of the elements of L does not divide k, then every L-intersecting family of subsets of $[n]$ has at most n members. (Comment: If $0 \in L$, $\gcd(L) \mid k$, and $k \ge |L|(\max(L))^2$, then when $n \ge 2k^2$ there is an L-intersecting family of size $(n/2k)^2$ in $\binom{[n]}{k}$.) (Alon–Babai [1980])

20.2.4. (◇) Let $X = \binom{t^3}{3}$. Color the complete graph with vertex set X by making an edge red if its endpoints have one common element and blue otherwise. Conclude for diagonal Ramsey numbers that $R(t,t) > \binom{t^3}{3}$. (Comment: The graph $(t-1)K_{t-1}$ gives $R(t,t) > (t-1)^2$.) (Nagy [1972])

20.2.5. An inner product space V is **nonsingular** if no nonzero vector is orthogonal to all of V. For a subspace U, let $U^\perp = \{v \in V \colon \langle v, u\rangle = 0 \text{ for all } u \in U\}$.

a) Prove that if V is a nonsingular n-dimensional inner product space, then $\dim U + \dim U^\perp = n$.

b) Prove that in Eventown (Example 20.2.1) at most $2^{\lfloor n/2\rfloor}$ clubs can be formed. That is, $[n]$ contains at most $2^{\lfloor n/2\rfloor}$ sets of even size such that the intersection of every pair of these sets has even size. (Berlekamp [1969])

20.2.6. By Theorem 20.2.5, in $\mathbb{R}^n$ there is no two-distance set with more than $(n+1)(n+4)/2$ points. Construct a two-distance set of size $\binom{n+1}{2}$ in $\mathbb{R}^n$. (Hint: Start with such a set in $\mathbb{R}^{n+1}$.)

20.2.7. (+) *Improved bound on the size of two-distance sets.*

a) In $\mathbb{R}^n$, the *affine hull* of vectors $v_1, \dots, v_m$ is $\{\sum_{i=1}^m c_i v_i \colon \sum_{i=1}^m c_i = 0\}$. Let B be the m-by-$(n+1)$ matrix whose ith row is v_i plus 1 in column $n+1$. Prove that if the affine hull is all of $\mathbb{R}^n$, then the columns of B are linearly independent.

b) Prove that if the columns of a real m-by-k matrix are linearly independent, then $B^T B$ is nonsingular.

c) Prove that when the polynomials $\{1, x_1, \dots, x_n\}$ are added to the set of polynomials constructed from a two-distance set in Theorem 20.2.5, the enlarged set of polynomials is linearly independent. Conclude that the maximum size of a two-distance set in $\mathbb{R}^n$ is at most $\binom{n+2}{2}$. (Blokhuis [1981])

20.2.8. Given $\mathcal{F}_1, \mathcal{F}_2 \subseteq \binom{[n]}{k}$, let $\sigma(\mathcal{F}_2)$ denote the permutation of $\mathcal{F}_2$ obtained by applying the permutation $\sigma\colon [n] \to [n]$.

a) Letting all permutations of $[n]$ be equally likely, prove that the expectation of $|\mathcal{F}_1 \cap \sigma(\mathcal{F}_2)|$ is $|\mathcal{F}_1||\mathcal{F}_2|/\binom{n}{k}$.

b) Suppose that S_1 and S_2 are disjoint, and suppose that $\mathcal{F}_1$ is S_1-intersecting and $\mathcal{F}_2$ is S_2-intersecting. Prove that $|\mathcal{F}_1| \cdot |\mathcal{F}_2| \le \binom{n}{k}$. (Szegedy [1990])

20.2.9. Choose $n, p \in \mathbb{N}$ with p prime and $n > 2p$. Let $G_{n,p}$ be the graph whose vertices are the incidence vectors of $(2p-1)$-sets in $[n]$, with two vertices adjacent when their distance in $\mathbb{R}^n$ is $\sqrt{2p}$. Prove that $\chi(G_{n,p}) \ge \binom{n}{2p-1}/\binom{n}{p-1}$. Improve the lower bound on the chromatic number of the unit-distance graph in $\mathbb{R}^n$ by choosing p to maximize the lower bound on $\chi(G_{n,p})$.

20.2.10. Use Stirling's Formula (Theorem 16.@.@) to approximate $\ln\binom{p^3}{p^2-1}/\ln\binom{p^3}{p-1}$ when p is large.

20.2.11. Prove that the points of $\mathbb{R}^2$ cannot be colored with three colors so that points at distance 1 have different colors. Prove that seven colors suffice. (Hint: For the upper bound, make use of a hexagonal grid.)

20.2.12. Prove that the points of $\mathbb{R}^2$ can be colored using $n^{n/2}$ colors so that points at distance 1 have different colors.

20.2.13. Assume Snevily's Theorem: If L is a set of s positive integers, then the size of an L-intersecting family of subsets of $[n]$ is at most $\sum_{i=0}^{s}\binom{n-1}{s}$. Prove Theorem 20.2.10 from this.

20.2.14. Show that Theorem 20.2.24 is a special case of Theorem 20.2.25. Prove Theorem 20.2.25. (Alon–Nathanson–Rusza [1996])

20.2.15. Let A and B be nonempty subsets of $\mathbb{Z}_p$, where p is prime. Prove that the number of sums $x+y$ such that $x \in A$, $y \in B$, and $xy \neq 1$ is at least $\min\{p, |A|+|B|-3\}$. (Hint: Use Theorem 20.2.25.) (Alon–Nathanson–Rusza [1995])

20.2.16. *More from Theorem 20.2.25.* (Alon–Nathanson–Rusza [1996])

a) Let $c_1, \dots, c_k$ be nonnegative integers summing to $m+\binom{k}{2}$, where $m \ge 0$. Prove that

$$\left[\prod_{i=1}^{k} x_i^{c_i}\right]\left(\sum_{i=1}^{k} x_i\right)^m \prod_{1\le i<j\le k}(x_j - x_i) = \frac{m!}{\prod_{i=1}^{k} c_i!}\prod_{1\le i<j\le k}(c_j - c_i).$$

(Hint: There is a proof using the Hook-length Formula (Theorem 16.2.@) and a more direct proof.)

b) Let $A_1, \dots, A_k$ be nonempty subsets of $\mathbb{Z}_p$, where p is prime. Let $S(A_1, \dots, A_k)$ denote the set of sums of distinct elements $a_1, \dots, a_k$ such that $a_i \in A_i$ for all i. Prove that if the sizes of $A_1, \dots, A_k$ are distinct and sum to less than $p+\binom{k+1}{2}$, then $|S(A_1, \dots, A_k)| > \sum_{i=1}^{k}|A_i| - \binom{k+1}{2}$. (Hint: Apply part (a) and Theorem 20.2.25.)

c) Let $A_1, \dots, A_k$ be nonempty subsets of $\mathbb{Z}_p$, where p is prime, indexed so that $|A_1| \ge \cdots \ge |A_k|$. Let $b_1 = |A_1|$, and let $b_i = \min\{b_{i-1}-1, |A_i|\}$ for $2 \le i \le k$. Prove that if $b_k > 0$, then $|S(A_1, \dots, A_k)| \ge \min\{p, \sum_{i=1}^{k} b_i - \binom{k+1}{2}+1\}$.

d) Conclude that if A is a nonempty subset of $\mathbb{Z}_p$, then the number of sums of s distinct elements of A is at least $\min\{p, s\,|A| - s^2 + 1\}$. (Comment: Theorem 20.2.24 is the special case $s = 2$.) (Dias da Silva–Hamidoune [1994])

20.2.17. *Cauchy–Davenport in higher dimensions.* The **Hopf–Stiefel function** relative to a prime p is a function $r \circ s$ of positive integers r and s, defined to be the smallest integer n such that $\binom{n}{k}$ is divisible by p whenever $n-r < k < s$. The value can be computed recursively from the base-p representations of r and s. Note that $r \circ s = r+s-1$ whenever $r+s \le p+1$.

a) Prove that if A and B are nonempty subsets of a finite vector space over $\mathbb{F}_p$, with $r = |A|$ and $s = |B|$, then $|A+B| \ge r \circ s$.

b) Prove that part (a) is sharp for all r and s. (Hint: In one dimension, part (a) reduces to the Cauchy–Davenport Theorem. Prove sharpness by generalizing the sharpness example for that theorem.) (Eliahou–Kervaire [1998])

20.2.18. Use the Combinatorial Nullstellensatz to prove that the minimum number of hyperplanes in $\mathbb{R}^n$ that do not contain 0 but together cover all the other points in $\{0,1\}^n$ is n. (Alon–Füredi [1993])

20.2.19. *The Permanent Lemma.*

a) Let A be an n-by-n matrix with nonzero permanent over a field $\mathbb{F}$. Prove that for any $b \in \mathbb{F}^n$ and sets $S_1, \dots, S_n$ of size 2 in $\mathbb{F}$, there is an vector $x \in \prod_{i=1}^{n} S_i$ such that Ax differs from b in every coordinate. (Alon–Tarsi [1989])

b) Let p be a prime. Prove that every list of $2p-1$ members of $\mathbb{Z}_p$ contains p terms that sum to 0 modulo p. (Erdős–Ginzburg–Ziv [1961])

20.2.20. Prove that the $1,2$-Conjecture holds for complete graphs. (Przybyło–Woźniak [2010])

20.2.21. Construct infinitely many trees that are not $(1,2)$-choosable. (Wong–Zhu [2009+])

20.2.22. ($\diamond$) Let G be a tree with an odd number of edges. Give each edge a list of size 2 and each vertex a list of size 1, except that one vertex has a list of size 2. Use Theorem 20.2.59 to prove that a proper total weighting of G can be chosen from the lists. Conclude that all trees are $(2,2)$-choosable. (Wong–Zhu [2009+])

20.2.23. ($\diamond$) For each vertex v in a graph G, specify $B(v) \in \{1, \dots, d_G(v)\}$.

a) Prove that if $\sum_{v\in V(G)}|B(v)| < |E(G)|$, then G has a nontrivial subgraph H such that $d_H(v) \notin B(v)$ for all $v \in V(G)$ (note that degree 0 is allowed, but not at all vertices). (Shirazi–Verstraëte [2008])

b) Show that part (a) is sharp by constructing infinitely many examples with $\sum_{v\in V(G)}|B(v)| = |E(G)|$ where no such subgraph exists.

20.2.24. Let p be an odd prime, and let k be an integer with $1 \le k < p$. Given $a_1, \dots, a_k \in \mathbb{F}_p$ and distinct elements $b_1, \dots, b_k \in \mathbb{F}_p$, prove that there is a permutation σ of $[k]$ such that the values $a_i + b_{\sigma(i)}$ are distinct modulo p. (Hint: Use the Vandermonde determinant in an application of the Combinatorial Nullstellensatz.) (Alon [2000])

20.2.25. Given a permutation σ of $[k]$, let $d_\sigma(i,j) = a-b$, where the positions of i and j in the word form are a and b, respectively.

a) For each pair (i,j) with $1 \le i < j \le k$, specify a forbidden distance $f(i,j)$. Prove that there is a permutation σ such that $d_\sigma(i,j) \neq f(i,j)$ for $1 \le i < j \le k$.

b) Let n and k be positive integers with $2k \le n+1$. Apply part (a) to prove that for any list $a_1, \dots, a_k$ of elements of $\mathbb{Z}_n$, there is a permutation $\sigma \in \mathbb{S}_k$ such that the elements $a_{\sigma(i)} + i$ for $1 \le i \le k$ are distinct modulo n.

(Comment: This result was used to show that every tree with n edges and radius r decomposes some complete graph with at most $32(2r+4)n^2+1$ vertices.) (Kézdy–Snevily [2002])

20.2.26. Let v be a vertex in a tree T. Prove that if a tree T' with m edges obtained by adding pendant edges at v to T, where $2m+1$ is prime and m is sufficiently large, then T' has a ρ-valuation. (Kézdy [2006])

20.2.27. Let T be a stunted tree with m edges in the canonical indexing $e_1, \dots, e_m$ used in Theorem 20.2.33, where $2m+1$ is prime. Modify the proof of Theorem 20.2.33 to show that if $S_1, \dots, S_m$ are subsets of $\mathbb{Z}_{2m+1}$ such that $|S_i| = i$ for $1 \le i \le m$, then T has a ρ-valuation such that the difference on edge e_i is not in S_i. (Kézdy [2006])