0.1 Intersection restrictions

0.1.1 When all intersections have the same size

In the Erdős-Ko-Rado theorem, both in its uniform and its not-necessarily-uniform versions, we forbid empty pairwise intersections. In other words the set of allowed sizes for pairwise intersections is $\{1, 2, \ldots, n-1\}$. What if we restrict us more and allow only a single intersection size? That is, let $\lambda \in \mathbb{N}$ be an integer and $\mathcal{F} \subseteq 2^{[n]}$ such that $|F_1 \cap F_2| = \lambda$ for every $F_1, F_2 \in \mathcal{F},\ F_1 \neq F_2$. How large can then $\mathcal{F}$ be?

We will tackle this problem in a bit, but we first collect ideas by considering a modular variant (with a story [?]).

Eventown vs. Oddtown.

In a little town called Eventown the 32 citizens love to organize various clubs. But they have to follow the strict traditional Eventown-rules of which there are two:

E1 Every club has to have an even number of members.

E2 Every pair of clubs has to have an even number of members in common.

The city council at some point is facing an administrative nightmare due to the number of clubs getting out of control. Indeed, if the citizens for example, were to pair themselves up and would join or not join clubs only together with their pairs, then the system of all $2^{16} = 65536$ such possibilities would create a feasible club system according to the above Eventown-rules. The mayor wants to cut down the number of clubs and considers changing the century old rules whose motivation is anyway lost in the obscurity of old times. After consultation with the wise, they consider the following slight modification of the Eventown-rules.

O1 Every club has to have an odd number of members.

O2 Every pair of clubs has to have an even number of members in common.

The difference is in only one word, the conditions **E2** and **O2** are the same. How many clubs could there be now? Let us first take a look at some construction ideas, with the set of citizens being denoted by $[n]$.

1) Taking the n singletons creates n clubs of size 1 each with all pairwise intersection 0.

2) When n is even, one could take the complement of singletons. This creates an $(n-1)$-uniform family with all pairwise intersections having size $n-2$.

3) Again for n even, we could consider the so-called *two-star* construction. Let $F_i = \{i, n-1, n\}$ for $1 \leq i \leq n-2$, $F_{n-1} = \{1, 2, ..., n-2, n-1\}$, and $F_n = \{1, 2, ..., n-2, n\}$. This system is not anymore uniform, but the possible set sizes are 3 and $n-1$, both odd and for the pairwise intersections we have $|F_i \cap F_j| \in \{2, n-2\}$ for $i \neq j$.

All these construction have n sets and there are in fact many more such constructions. The next theorem shows that one cannot do better.

Theorem 0.1.1 (Oddtown Theorem, Berlekamp, 1969)**.** *Let $\mathcal{F} \subseteq 2^{[n]}$ be a set family satisfying both* **O1** *and* **O2***. Then $|\mathcal{F}| \leq n$.*

This theorem is pretty significant for the mayor of Eventown. With changing just a single word in the rules, the number of possible clubs is reduced from exponential to linear. The council votes to change the name of the town to Oddtown and they live happily ever after.

Before the actual proof of the Oddtown Theorem we make a few general comments in order to motivate and introduce the general method that will be of use in this chapter.

In a typical extremal combinatorial problem, the greater the number of extremal families[1], the less likely that a purely combinatorial argument will lead to a solution, since a proof eventually must consider all extremal structures, and be tight for each of them in each of the proof steps. If these families are combinatorially very different, this might necessarily lead to an unmanageable number of combinatorial case distinctions.

For some of these problems the stars are aligned and the difficulties posed by multiple extremal examples can be mitigated by realizing that the combinatorial problem, or rather its extremal structures, hides the features and concepts of another mathematical discipline in the background. In such cases, the simplest, or most efficient descriptions of extremal structures are not necessarily combinatorial, but might have to be formulated in another language, which could be algebraic, probabilistic, or, even topological. Solutions of this sort, connecting different branches of mathematics, are considered gems: they are rare and beautiful.

The Oddtown Theorem is one of those situations, where the the extremal families are far from being unique. In fact one can prove that their number is super-exponential [**?**][Exercise 1.1.14]. For the (easy) proof of Theorem 0.1.1 one only has to realize that the right language of the problem is the one of linear algebra. And then even though the number of extremal set-systems is superexponential and the feasibility of their combinatorial characterization is questionable at best, they have a very simple linear algebraic description as the orthonormal bases in $\mathbb{F}_2^n$.

The connection between combinatorics and algebra is provided through the characteristic vector $\mathbf{v}_F \in \{0,1\}^n$ of sets $F \subseteq [n]$, where $(v_F)_i = 1$ if and only if $i \in F$.

The key realization relevant to families with pairwise intersection restrictions is that the size of the intersection of two sets is equal to the standard inner product of the two characteristic vectors:

$$|A \cap B| = \sum_{i=1}^{n} (\mathbf{v}_A)_i (\mathbf{v}_B)_i =: \mathbf{v}_A \cdot \mathbf{v}_B.$$

Indeed, when calculating a term $(\mathbf{v}_A)_i(\mathbf{v}_B)_i$ of the sum, we have a 1 if and only if $i \in A \cap B$.

This connection will allow us to translate the combinatorial condition we have on a family into linear algebra and use this information to derive the linear independence of the characteristic vector. This in turns makes the dimension of the space an upper bound on their number.

The simple linear algebra fact that the size of a linearly independent set of vectors is *at most* the dimension of the ambient vector space is called the *dimension bound.* It is the simplest manifestation of the Linear Algebra method.

Proof of the Oddtown Theorem. Let $\mathcal{F} = \{C_i : 1 \le i \le m\}$. To each set C_i we associate its characteristic vector $\mathbf{v}_i \in \{0,1\}^n$. From rules **O1** and **O2**, using $\mathbf{v}_i \cdot \mathbf{v}_j = |C_i \cap C_j|$, we infer that $\mathbf{v}_i \cdot \mathbf{v}_i$ is odd for every $i \in [n]$ and $\mathbf{v}_i \cdot \mathbf{v}_j$ is even for every $i \neq j$.

We would like to claim the linear independence of the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_m$, for which should name a field to work over. Considering our conditions on the pairwise dot products, it does not come as a great shock that we choose to show linear independence over the two-element field $\mathbb{F}_2$.

Suppose we have a linear combination $\sum_{i=1}^{m} \alpha_i \mathbf{v}_i = \mathbf{0}$, with $\alpha_i \in \mathbb{F}_2$. For every $j \in [m]$, we take the

[1] In an extremal combinatorial problem an *extremal family* is one with the largest (smallest) possible number of edges among those with the required property. In our problem these are the set families of size n that satisfy both **O1** and **O2**.

dot product of it with $\mathbf{v}_j$ and obtain

$$0 = \mathbf{0} \cdot \mathbf{v}_j = \left(\sum_{i=1}^{m} \alpha_i \mathbf{v}_i \right) \cdot \mathbf{v}_j = \sum_{i=1}^{m} \alpha_i \left(\mathbf{v}_i \cdot \mathbf{v}_j \right) = \alpha_j \left(\mathbf{v}_j \cdot \mathbf{v}_j \right) = \alpha_j .$$

Here we used that all but one of the dot products are zero over $\mathbb{F}_2$. This implies that $\alpha_j = 0$ for every j, consequently the vectors are indeed linear independent. Hence, their number cannot be more than the dimension of the space and we get $m \leq \dim \mathbb{F}_2^n = n$. □

Fisher's Inequality

Let us return to our original problem, where we restricted the size of every pairwise intersection to a unique non-zero integer λ. It turns out that the same upper bound holds here as well.

Theorem 0.1.2 (Non-uniform Fisher Inequality)**.** *Let $n \in \mathbb{N}$ and suppose $1 \leq \lambda \leq n$. If $\mathcal{F} \subseteq 2^{[n]}$ satisfies $|F_1 \cap F_2| = \lambda$ for every $F_1, F_2 \in \mathcal{F}$, $F_1 \neq F_2$, then $|\mathcal{F}| \leq n$.*

Note that if $\lambda = 0$, then the sets F_i must be pairwise disjoint and hence we can have at most $n+1$ such sets (including the $\emptyset$).

Fischer's Inequality originates from statistics/experiment design, where the problem is how to divide test subject between treatments fairly. There are several other versions of the inequality, but this one is more general.

Constructions: for $\lambda = 1$ we can take near pencils (see Figure 1) and (nondegenerate) projective planes. The case $\lambda = 1$ is the De Bruijn-Erdős theorem. For $\lambda = n-2$ we can take the complements of singletons.

It is an open problem to classify all cases of equality. Can we have equality for all λ?

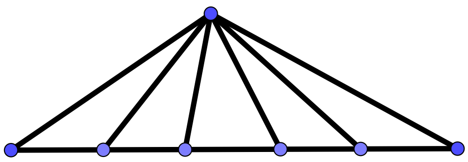

Figure 1: A near pencil, a.k.a., degenerate projective plane

In the proof we plan to use the same idea we developed for the Oddtown Theorem: we show that the characteristic vectors $\mathbf{v}_F$ of a family $\mathcal{F}$ from the theorem is an independent set of vectors. We have liberty in choosing the field we want work over, since $0, 1$-vectors make sense over any field. The right choice here will be a field with zero characteristic, say $\mathbb{R}$ or $\mathbb{Q}$.

In the Oddtown Theorem we inferred linear independence directly by considering a linear combination equl to $\mathbf{0}$ and showing it is trivial. An alternative view of linear independence is through the Gram-matrix $M(V(\mathcal{F}))$ of the characteriztis vectors.

Definition 0.1.3. *The Gram-matrix $M(V)$ of a set of vectors $V = \{\mathbf{v}_1, \ldots, \mathbf{v}_m\} \subseteq \mathbb{F}^n$, over some field $\mathbb{F}$, is the $|V| \times |V|$-matrix with entries*

$$m_{i,j} = \mathbf{v}_i \cdot \mathbf{v}_j .$$

Note that the Gram-matrix is equal to AA^T, where A is the $m \times n$-matrix whose ith row is $\mathbf{v}_i$. Therefore $\mathrm{r}k(M) \leq \mathrm{r}k(A) \leq m$.[2] So if the Gram matrix is non-singular over any field then the vectors are linear independent over that field.

Proposition 0.1.4. *If $rk(M) = m$ over some field $\mathbb{F}$ then $m \leq n$.*

[2]For any two matrices $B \in \mathbb{F}^{k \times \ell}$ and $C \in \mathbb{F}^{\ell \times m}$ we have that $\mathrm{r}k(BC) \leq min\{rk(B), \mathrm{r}k(C)\}$, since the column space of BC is contained in the column space of C (the dimension of which is $\mathrm{r}k(C)$ and the row space of BC is contained in the row space of C (the dimension of which is $\mathrm{r}k(C)$).

The Gram matrix as a real matrix is positive semidefinite. Indeed, for every $x \in \mathbb{F}^m$ we have

$$xMx^T = xAA^Tx^T = (xA)(xA)^T = ||xA||^2 \geq 0.$$

Considering the Gram matrix as a real matrix and checking whether it is positive definite gives an alternative way to check whether the vectors are linear independent.
Proposition 0.1.5. *A set $V \subseteq \mathbb{R}^n$ of real vectors is linearly independent if and only if their Gram-matrix $M = M(V)$ is positive definite.*

Proof. The vector set V is linearly independent if and only if for every $x \in \mathbb{F}^m \setminus \{\mathbf{0}\}$, the linear combination xA of the rows of A is non-zero. This happens if and only if $xMx^T = ||xA||^2 > 0$ for every $x \in \mathbb{F}^m \setminus \{\mathbf{0}\}$, that is if M is positive definite. □

Proof. Let $\mathcal{F} = \{F_1, ..., F_m\}$ be a family, such that $|F_i \cap F_j| = \lambda$ for every $i \neq j$.

Case 1. There is set $F_i \in \mathcal{F}$ of size $\lambda = |F_i|$. Then, $|F_i \cap F_j| = \lambda$ implies that $F_i \subseteq F_j$ for every j. Since $F_j \cap F_\ell \supseteq F_i$ is also of size λ, the sets $\{F_j \setminus F_i\}_{j=1}^m$ must be pairwise disjoint and are contained in $[n] \setminus F_i$. Hence $m \leq n - \lambda + 1 \leq n$.

Case 2. We have $|F_i| > \lambda$ for every i. Let us introduce the positive integers $\gamma_i = |F_i| - \lambda > 0$. Let $\mathbf{v_i}$ be the characteristic vector of F_i for $i \in [m]$. We claim that this set of m vectors are linearly independent in $\mathbb{R}^n$. Indeed, in this case the Gram matrix is positive definite, since

$$xMx^T = \sum_{i,j} x_i x_j m_{ij} = \sum_{i,j} x_i x_j \lambda + \sum_{i=1}^m x_i^2 \gamma_i = \left(\sum_{i=1}^m x_i \sqrt{\lambda} \right)^2 + \sum_{i=1}^m x_i^2 \gamma_i > 0$$

for every $x \in \mathbb{F}^m \setminus \{\mathbf{0}\}$. Therefore the vectors of $V(\mathcal{F})$ are linearly independent, so their number is not more than the dimesnion of the space, which is n.

□

Remark 0.1.6. *Note then when talking about linear independence, one has to consider the field behind the ambient vector space. Vectors that independent over $\mathbb{R}$ might be dependent over $\mathbb{F}_2$ (the converse however is not possible!).*

0.1.2 Multiple permissable intersection sizes

Restricting the *number* of possible intersections

Problem/Result	Set size	Intersection size	Bound on family size	Total # of sets
Intersecting	any	$\neq 0$	2^{n-1}	2^n
Erdős-Ko-Rado	$k\,(\leq n/2)$	$\neq 0$	$\binom{n-1}{k-1}$	$\binom{n}{k}$
Eventown	$0 \pmod 2$	$0 \pmod 2$	$\geq 2^{\lfloor n/2 \rfloor}$	2^{n-1}
Oddtown	$1 \pmod 2$	$0 \pmod 2$	n	2^{n-1}
Fisher's Ineq.	any	$= \lambda \quad (1 \leq \lambda \leq n)$	n	2^n

Definition 0.1.7. *Let L be a set of non-negative integers. The family $\mathcal{F}$ is* L-intersecting *if*

$$|F_1 \cap F_2| \in L \qquad \forall\ F_1, F_2 \in \mathcal{F},\ F_1 \neq F_2.$$

The family $\mathcal{F}$ is $\pmod p$-L-intersecting *if*

$$|F_1 \cap F_2| \in L \pmod p \quad \forall F_1, F_2 \in \mathcal{F}, F_1 \neq F_2.$$

The following fundamental question about L-intersecting families is largely open.

Question 0.1.8. *Given some set L, how large can an L-intersecting/(mod p)-L-intersecting (k-uniform) family $\mathcal{F} \subseteq 2^{[n]}$ be?*

In the previous sections we settled several cases.

Examples:

- If $L = \{1, 2, \ldots, n\}$ then the largest L-intersecting family is of size 2^{n-1}.
- If $L = \{1, 2, \ldots, n\}$ then the largest k-uniform L-intersecting family is $\binom{n-1}{k-1}$.
- $L = \{\lambda\}$, then the largest L-intersecting family is of size at most n provided $1 \leq \lambda \leq n$ (and equal to $n+1$ provided $\lambda = 0$).
- $L = \{0\}$, then the largest (*mod* 2)-L-intersecting family is of size at least $2^{\lfloor \frac{n}{2} \rfloor}$.
- $L = \{0\}$, then the largest (*mod* 2)-L-intersecting (*mod* 2)-1-uniform family is of size n.

While no-one knows a complete answer to Question 0.1.8, we can provide sharp answers that depend *only* on the cardinality of L. In the previous section we dealt with cases when $|L| = 1$.

First we consider the modular case, i.e. extensions of the Oddtown Theorem. To this end we introduce a far-reaching new point of view to the linear independence of vectors, involving polynomials. The obtained General Mod-p-Town Theorem will also facilitate unexpected applications in geometry. In order to tackle the uniform case and non-modular problems, in particular the generalizations of Fischer's Inequality, we need to add a couple of useful enhancements to the existing toolbox.

Modular intersection restrictions—Extensions of the Oddtown Theorem

From the proofs of the Oddtown Theorem and the Fischer Inequality it is clear that whenever we find a field over which the intersection matrix is non-singular then the size of the family is bounded by the number of points. In this direction, we state explicitely the following straightforward generalization of the Oddtown Theorem.

Theorem 0.1.9 (Mod-p-Town Theorem)**.** *Let p be a prime and let $\mathcal{F} \subseteq [n]$ be a (mod p)-$\{0\}$-intersecting family such that $|F| \not\equiv 0 \pmod p$ for $\forall F \in \mathcal{F}$. Then*

$$|\mathcal{F}| \leq n.$$

Proof. Considering the Gram matrix of the characteristic vectors of such a family $\mathcal{F}$, we find that it is a diagonal matrix over the p-element field $\mathbb{F}_p$, with diagonal elements that are not 0. Hence the Gram matrix is non-singular over $\mathbb{F}_p$ and in turn the characteristic vectors are linearly independent and their number is at most the dimension of the space. □

If we omitted the restriction on the sizes of the sets, one can construct (analogously to the Eventown construction) a (*mod* p)-$\{0\}$-intersecting family of size $2^{\lfloor n/p \rfloor}$. It turns out that this restriction, forbidding set sizes to be equal to an intersection size $\pmod p$, is enough to guarantee a tight bound that is polynomial in n (as opposed to exponential), no matter how many elements L has. The degree of the polynomial depends on $|L|$ though.

Theorem 0.1.10 (General Mod-p-Town Theorem)**.** *Let p be a prime, L be a set of s integers, and $\mathcal{F} \subseteq 2^{[n]}$ be a* $\pmod p$*-L-intersecting family.*

(a) [Deza-Frankl-Singhi, 1983] If $|F| \notin L \pmod p$ for every $F \in \mathcal{F}$, then

$$|\mathcal{F}| \leq \binom{n}{s} + \binom{n}{s-1} + \cdots + \binom{n}{0};$$

(b) *[Frankl-Wilson, 1981] If $\mathcal{F}$ is $\pmod p$-k-uniform (that is $|F| \equiv k \pmod p$ for every $F \in \mathcal{F}$), for some $k \notin L \cup [0, s-1] \pmod p$ then*

$$|\mathcal{F}| \leq \binom{n}{s}.$$

Remark. Theorem 0.1.10 part (b) is tight and part (a) is asymptotically tight, for any prime p and positive integer $s < p$. Indeed, the family $\mathcal{F} = \binom{[n]}{s}$ is $\pmod p$-L-intersecting with $L = \{0, 1, \ldots, s-1\}$ and $|F| = s \notin L \pmod p$ for every $F \in \mathcal{F}$.

The main point of the proof of the General Mod-p-Town Theorem is how to adapt the dimension argument of the Oddtown Theorem to this more general set-up. In our original framework we associated with each set F_i its characteristic vector $\mathbf{v}_i \in \{0,1\}^n$ and proved that they are linearly independent. This cannot anymore work as the upper bound we seek[3] is much larger than the dimension n of the space these vectors live in.

Observe however that the linear independence of vectors $\mathbf{v}_i \in \mathbb{F}^n$ is equivalent to the linear independence of the corresponding linear functions $\ell_i(\mathbf{x}) = (\mathbf{v_i})_1 x_1 + \ldots + (\mathbf{v_i})_n x_n$. This motivates us to associate with each set $F_i \in \mathcal{F}$ a *function* instead of a vector, and prove the linear independence of these functions.

To this end the following criterion will be convenient.

Lemma 0.1.11. *Let $\mathbb{F}$ be a field and let $f_1, \ldots, f_m \in \mathbb{F}[x_1, \ldots x_n]$ polynomials in n variables. If there exist vectors $\mathbf{v_1}, \ldots, \mathbf{v_m} \in \mathbb{F}^n$ such that*

- *$f_i(\mathbf{v_i}) \neq 0$ for all $i \in [m]$, and*
- *$f_i(\mathbf{v_j}) = 0$ for all $j \neq i$,*

then $f_1, \ldots, f_m$ are linearly independent.

Proof. Suppose $\lambda_1 f_1 + \cdots + \lambda_m f_m = 0$. Substituting $\mathbf{v_i}$, the equality reduces to $\lambda_i f_i(\mathbf{v_i}) = 0$. Since $f_i(\mathbf{v_i}) \neq 0$, we infer $\lambda_i = 0$. □

Proof of Theoerm 0.1.10. In the Mod-p-Town Theorem the intersection conditions $|F_i \cap F_j| \equiv 0 \pmod p$, $i \neq j$, translated to the orthogonality of the characteristic vectors over $\mathbb{F}_p$, which (together with them *not* being self-orthogonal) took care of their linear independence. Now, instead of $\mathbf{v}_i \cdot \mathbf{v}_j = 0$, we might have $\mathbf{v}_i \cdot \mathbf{v}_j = \ell$ for various different $\ell \in L \subseteq \mathbb{F}_p$. To encode the intersection conditions consider the product

$$\prod_{\ell \in L} (\mathbf{v}_i \cdot \mathbf{v}_j - \ell) = \prod_{\ell \in L} (|F_i \cap F_j| - \ell)\,.$$

By the conditions of our theorem this product is 0 for every $i \neq j$ and non-zero for $i = j$.

To each F_i let us associate the polynomial function $f_i : \mathbb{F}_p^n \to \mathbb{F}_p$ defined by

$$f_i(\mathbf{x}) = \prod_{\ell \in L} (\mathbf{v}_i \cdot \mathbf{x} - \ell) = \prod_{\ell \in L} ((\mathbf{v_i})_1 x_1 + \cdots + (\mathbf{v_i})_m x_m - \ell)\,.$$

Then the conditions of our theorem translate exactly to the conditions of Lemma 0.1.11 and hence the functions $f_1, \ldots, f_m$ are linearly independent. Therefore their number cannot be more than the dimension of the $\mathbb{F}_p$-vector space $\mathbb{F}_p[x_1, \ldots x_n]$. This is waaay too large for our purposes, but fortunately the functions $f_1, \ldots, f_m$ live in a much smaller space: the subspace they generate.

What is then this dimension $\dim(\langle f_1, \ldots, f_m \rangle)$? Each f_i is the product of s linear functions in the n variables $x_1, \ldots, x_n$. Expanding the parenthesis, each f_i is the linear combination of the

[3] that we actually *need* to seek, due to the large construction in the Remark

monomial terms of the form $x_1^{s_1} \cdot \dots \cdot x_n^{s_n}$ with $s_1 + \dots + s_n \le s$. The number of these terms is equal to $|\{(s_0, s_1, \dots, s_n) \in \mathbb{N}_0^{n+1} : \sum_{i=0}^n s_i = s\}|$, i.e. the number of ways to distribute s indistinguishable balls into $n+1$ labeled boxes: $\binom{n+s}{s}$. This is still too big compared to what we set out to show.

Multilinearization. We introduce yet another trick, called *multilinearization* to reduce the dimension. We will make use of the fact that the characteristic vectors (witnessing the linear independence in the Lemma) have only 0 or 1 coordinates.

From f_i define $\tilde{f}_i$ by expanding the product and replacing each power x_i^k by a term x_i for every $k \ge 1$ and i, $1 \le i \le m$. So for example, the term $x_2^3 x_5 x_9^{11} x_{10}^2$ is replaced by $x_2 x_5 x_9 x_{10}$.

Since $0^k = 0$ and $1^k = 1$ for every $k \ge 1$ we have that $f_i(\mathbf{v_j}) = \tilde{f}_i(\mathbf{v_j})$ for every i, j. Consequently the conditions of the lemma remain valid for the $\tilde{f}_i$s and the (now) multi**linear** polynomials $\tilde{f}_1, \dots, \tilde{f}_m$ of total degree s are also linearly independent.

Bounding the dimension for part (a). These polynomials live in a space spanned by the basic monomials $\prod_{j=1}^k x_{i_j}$ of degree at most s. Such monomials are determined by a subset of the variables of size at most s, hence their number is equal to

$$\binom{n}{s} + \binom{n}{s-1} + \dots + \binom{n}{1} + \binom{n}{0}.$$

Mod-p uniform families. For part (b) we need to strengthen the bound using the fact that all members of the family have the same cardinality k modulo p. We do this by adding more functions to the family $\tilde{f}_1, \dots, \tilde{f}_m$, that are also part of the vector space of multilinear polynomials of degree at most s over $\mathbb{F}_p$, and show that this larger set of functions is still linearly independent!

A natural choice is the set of monomials $\prod_{i \in I} x_i$ corresponding to the subsets $I \subseteq [n]$ of size at most $s-1$. These polynomials are part of the standard basis of $\mathbb{F}_p[x_1, \dots x_n]$ and hence are linearly independent. In order to show their linear indepedence from $\tilde{f}_1, \dots, \tilde{f}_m$ as well, they are multiplied with the factor $(\sum_{i=1}^n x_i - k)$.

Formally, for $|I| \le s-1$, let $g_I(\mathbf{x})$ be the multilinearized version of $(\sum x_i - k) \prod_{i \in I} x_i$. Suppose

$$\sum_{i=1}^n \lambda_i \tilde{f}_i + \sum_{|I| \le s-1} \mu_I g_I \equiv 0 \tag{1}$$

in $\mathbb{F}_p[x_1, \dots x_n]$. When substituting the characteristic vector $\mathbf{v_i}$ into this linear combination, what we are left with is $\lambda_i \tilde{f}_i(\mathbf{v_i}) = 0$, which implies $\lambda_i = 0$ (since $\tilde{f}_i(\mathbf{v_i}) = \prod_{\ell \in L}(k - \ell) \neq 0$, as $k \notin L$ (mod p)). Indeed $\tilde{f}_j(\mathbf{v}_i) = 0$ for $j \neq i$ (since pairwise intersection sizes are in L) and $g_I(\mathbf{v}_i) = 0$ (since $\sum_{j=1}^n (\mathbf{v_i})_j - k = 0$ in $\mathbb{F}_p$, due to $\mathcal{F}$ being (mod p)-k-uniform).

Once we know that all λ_i are 0 , we can show that each μ_I is 0 the standard way: substituting the characteristic vectors $\mathbf{v_J}$ of the subsets $J \subseteq [n]$ of size at most $s-1$. Formally, define a total ordering $\prec$ on the subsets of $[n]$ with cardinality $\le s-1$ that preserves the subset relation. Supposing that (1) is non-trivial, let J be the minimal subset according to $\prec$ for which $\mu_J \neq 0$. Substituting $\mathbf{v}_J$ into (1) we have that

- for $I \prec J$, $\mu_I g_I(\mathbf{v}_J) = 0$, since $\mu_I = 0$ by the minimality of J;
- For $I \succ J$, $\mu_I g_I(\mathbf{v}_J) = 0$, since $I \not\subseteq J$ implies $\prod_{i \in I} (\mathbf{v}_J)_i = 0$;
- For $I = J$ we have $\mu_J g_J(\mathbf{v}_J) = \mu_J \prod_{i \in J} (\mathbf{v}_J)_i (\sum_{i=1}^n (\mathbf{v_J})_i - k) = \mu_J \cdot 1 \cdot (|J| - k) \neq 0$ in $\mathbb{F}_p$, since $k \notin [0, s-1]$ (mod p) and $\mu_J \neq 0$ by the definition of J.

This contradicts that the substitution of $\mathbf{v_J}$ into (1) should be 0. Hence the linear combination is trivial and the functions involved are linearly independent.

So we found a set of $m + \sum_{i=0}^{s-1} \binom{n}{i}$ linearly independent functions in the space of multilinear polynomials of degree s in n variables. The dimension of this space is $\sum_{i=1}^{s} \binom{n}{i}$, implying the promised bound on m. □

Remark Note that the upper bound in part (a) of the above theorem can be reduced by 1 if $0 \in L$. Indeed, in that case the constant monomial 1 is *not* part of the span of $\tilde{f}_1, \dots, \tilde{f}_m$.

Remark. The theorem also holds if instead of requiring that k is not among $0, 1, \dots, s-1$ modulo p, we assume only that $k + s \leq n$. The same set of polynomials will still be linearly independent, but we need to use Inclusion-Exclusion to work around that now $\sum(\mathbf{v_J})_i - k$ *could be* 0 for some $|J| \leq s-1$.

Non-modular intersection theorems—Extentions of Fischer's Inequality

Our next theorem extends Fischer's Inequality and is tight for every value of the parameters. Note that, unlike in the General Mod-p-Town Theorem, we do not need to require anything about the sizes of the sets of the family.

Theorem 0.1.12. *Let L be a set of s integers and let $\mathcal{F}$ be an L-intersecting family.*

(a) [Frankl-Wilson, 1981] Then $|\mathcal{F}| \leq \sum_{i=0}^{s} \binom{n}{i}$.

(b) [Ray-Chaudhuri–Wilson, 1969] If additionally $\mathcal{F}$ is uniform then $|\mathcal{F}| \leq \binom{n}{s}$.

Proof. First note that part (b) readily follows from part (b) of the General Mod-p-Town Theorem if we use it with a prime $p > n$. Indeed, then an L-interscting family on vertex set $[n]$ is also $\pmod p$-L-intersecting and a k-uniform family is also $\pmod p$-k-uniform.

For part (a) let us first consider the same functions $f_i(\mathbf{x}) = \prod_{\ell \in L}(\mathbf{v}_i \cdot \mathbf{x} - \ell)$ that we used in the General Mod-p-Town Theorem. It will still be true that $f_i(\mathbf{v_j}) = 0$ whenever $i \neq j$, since $\mathbf{v_i} \cdot \mathbf{v_j} = |F_i \cap F_j|$ is equal to some element of L. For the diagonal entries however we might have $f_i(\mathbf{v_i}) = 0$, since nothing forbids now that a cardinality $|F_i|$ is an element of L. Consequently we cannot anymore use the diagonal criterion of Lemma 0.1.11.

To avoid vanishing diagonal entries we modify our functions and define $f_i : \mathbb{R}^n \to \mathbb{R}$, by taking the product of only those factors that correspond to elements of L which signal a *proper* intersection with F_i. That is,

$$f_i(\mathbf{x}) = \prod_{\ell \in L, \ell < |F_i|} (\mathbf{x} \cdot \mathbf{v}_i - \ell).$$

This way the diagonal entries are alright: $f_i(\mathbf{v_i}) \neq 0$ for every $i = 1, 2, \dots, m$. On the other hand some of the non-diagonal entries now might become non-zero, so Lemma 0.1.11 cannot be used again. However, this problem can be circumvented by arranging the sets in an order, so that the matrix of substitutions is upper triangular.

The following generalization of Lemma 0.1.11 states that linear independence can be concluded when the matrix of substitutions is non-singular.

Lemma 0.1.13. *Let $\mathbb{F}$ be a field and let $f_1, \dots, f_m \in \mathbb{F}[x_1, \dots x_n]$ polynomials in n variables. If there exist vectors $\mathbf{v_1}, \dots, \mathbf{v_m} \in \mathbb{F}^n$ such that the matrix $M = \left(f_i\left(\mathbf{v_j}\right)\right)_{i,j}$ of substitutions is non-singular, then $f_1, \dots, f_m$ are linearly independent.*

Proof. Suppose there is a vanishing linear combination $\lambda_1 f_1 + \cdots + \lambda_m f_m = 0$. Then the jth entry of the vector $(\lambda_1, \dots \lambda_n)M$ is equal to the substitution of $\mathbf{v_j}$ into the linear combination. Consequently $(\lambda_1, \dots \lambda_n)M = \mathbf{0}$. Since M is non-singular, we infer that all $\lambda_i = 0$. □

Let us order the elements of $\mathcal{F}$ such that $|F_1| \leq |F_2| \leq \cdots \leq |F_m|$. Then we have

$$f_i(\mathbf{v_j}) \begin{cases} \neq 0 & \text{if } i = j; \\ = 0 & \text{if } i > j. \end{cases}$$

Indeed, if $i > j$ then $\mathbf{v_i} \cdot \mathbf{v_j} = |F_i \cap F_j| < |F_i|$, since $|F_j| \leq |F_i|$. This means that the substitution matrix is upper-triangular with non-zero diagonal entries. Lemma 0.1.13 applies and $f_1, \ldots, f_m$ are linearly independent.

Now we can finish exactly as we did in the proof of the General Mod-p-Town Theorem: using the multilinear versions $\tilde{f}_i$ instead of f_i. Since the vectors witnessing the linear independence of the f_is in Lemma 0.1.13 are 0/1-vectors, the substituted values do not change and the multilinearized versions $\tilde{f}_i$ are also independent. They are again multilinear polynomials of degree at most s in n variables, hence the dimension count does not change. $\square$

Remark. 1. The theorem is tight for every s. For $L = \{0, 1, 2, \ldots, s-1\}$ the family $\mathcal{F} = \{F \subseteq [n] : |F| \leq s\}$ is L-intersecting.

2. Part (b) was the first in the row of these type of intersection theorems and an important step in the development of the linear algebra method. The modular versions were developed later by Frankl and Wilson in order to give better explicit construction of Ransey graphs. The General Mod$p$-Town Theorem was later also used to various geometric problems. A couple of these applications will be discussed in the next section.

0.1.3 Applications in Geometry

There are numerous applications of intersection theorems in geometry. The connection, just like the one to linear algebra, is quite simple. Interpreting characteristic vectors as points in the euclidean space $\mathbb{R}^n$, we obtain that the distance between $\mathbf{v}_A$ and $\mathbf{v}_B$ is exactly the size of the symmetric difference of A and B. In case $|A| = |B| = k$, we have that the distance

$$\|\mathbf{v}_A - \mathbf{v}_B\| = \sqrt{|(A \setminus B) \cup (B \setminus A)|} = \sqrt{2k - 2|A \cap B|}. \tag{2}$$

Hence restricting pairwise intersection sizes within a k-uniform family means restricting the distances between the corresponding points. This idea was applied succesfully in several long-standing, classic open problem in geometry.

Chromatic number of the unit distance graph

A fascinating open problem in combinatorial geometry asks to color the points of the Euclidean plane, with as few colors as possible, such that no pairs of points at distance one have the same color.

Formally, the unit distance graph UD_n of the n-dimensional euclidean space is defined on the vertex set $V(UD_n) = \mathbb{R}^n$ with edge set

$$E(G_n) = \{\mathbf{xy} : \|\mathbf{x} - \mathbf{y}\| = 1\}.$$

Problem 0.1.14 (Hadwiger-Nelson problem)**.** *What is the chromatic number $\chi(UD_2)$ of the plane?*

The best known bounds still allow the possibility of four different values.
Proposition 0.1.15. $4 \leq \chi(UD_2) \leq 7$.

Proof. For the lower bound we can find a subgraph of UD_2 with chromatic number 4. (The smallest such example is on 7 vertices.) For the upper bound one can give an explicit 7-coloring of the plane without a monochromatic pair of points at distance one using a tiling of the plane into regular hexagons with diameter slightly less than 1. □

The problem has interesting ties to the axiomatization of set theory. De Bruijn and Erdxős, using the Axiom of Choice, showed that the chromatic number is attained on a finite subset of the plane. Soifer and Shelah present evidence that the answer might depend on the ste of axioms we choose for set theory. It is also known that th chromatic numer is at least 5 if all color classes should be measurable and it is at least 6 if all color classes are the unions of faces of a locally finite plane graph.

The value of the chromatic number of the 3-space is known to be between 6 and 15. The value for large n is also an old question.
Problem 0.1.16. *How fast does $\chi(G_n)$ grow?*

Since we are talking about the chromatic number of an infinite graph, even finiteness is a question at first.
Homework 0.1.17. $\chi(UD_n) \leq n^{n/2} \cdot 2^n$.

An immediate lower bound of $\chi(UD_n) \geq n+1$ is given by a regular simplex of side length 1, which hosts a subgraph of UD_n isomorphic to K_{n+1}.

The superexponential upper bound was improved by Larman and Rogers (1972) to exponential ($\chi(UD_n) \leq const^n$ (HW)) and the linear lower bound to quadratic ($\chi(UD_n) = \Omega(n^2)$).

The huge gap between the bounds remained until Frankl and Wilson managed to find a way to bring extremal set theory to the rescue and improve the lower bound to exponential.
Theorem 0.1.18 (Frankl-Wilson, 1981)**.** $\chi(G_n) \geq \Omega(1.1^n)$.

Proof. Similar to the above lower bound on $\chi(UD_2)$, our lower bound on $\chi(UD_n)$ will be given by estimating the chromatic number of a well-chosen finite subgraph $H_n \subseteq UD_n$. To that end we plan to upper bound the independence number of H_n.

For H_n we use some of the vertices of the hypercube $\{0,1\}^n \subseteq \mathbb{R}^n$, which we interpret as characteristic vectors of a k-uniform set family $\mathcal{H}_n \subseteq 2^{[n]}$.

One small problem with this idea is that there is no unit-distance within the set $B_{k,n} = \{\mathbf{v} \in \{0,1\}^n : \sum \mathbf{v}_i = k\}$, so UD_n restricted to $B_{k,n}$ has chromatic number 1. This is easily overcome by realizing that unit-distance plays no special role in our problem. The unit distance graph UD_n and the "δ-distance graph" UD_n^δ (with vertex set $V(UD_n^\delta) = \mathbb{R}^n$ and edge set $E(UD_n^\delta) = \{\mathbf{xy} : \|\mathbf{x}-\mathbf{y}\| = \delta\}$) are isomorphic. So we could be flexible with which particular distance we care for, and will choose one for which $\alpha(UD_n^\delta)$ is small.

Any subset I of $V(UD_n^\delta)$ is independent if no pairs of points have a distance δ between hem. For a subset $I \subseteq V(H_n)$ this can be reformulated in the language of set families: the corresponding subfamily $\mathcal{I} \subseteq \mathcal{H}$ should not have two sets with a pairwise intersection of size $k - \frac{\delta^2}{2} =: \ell$. In other words the corresponding family $\mathcal{I}$ should be $L = [0, k-1] \setminus \{\ell\}$-intersecting. Our intersection theorems give an upper bound $\sum_{i=0}^{k-1} \binom{n}{i}$. This is not so good as it is only a factor n away from $\binom{n}{k}$ which is the number of vertices.

To obtain a stronger upper bound on the independence number we plan to use the modular intersection theorem instead with some prime p. To make this work we need that ℓ is the only representative of its $\pmod p$-residue class below k, that is if $\ell + p \geq k$. Then a k-uniform L-intersecting family is also $\pmod p$-L'-intersecting with $L' = [0, p-1] \setminus \{\ell\}$. To use the General Mod-p-town Theorem we also must have that $k \notin L' \pmod p$, that is $k \equiv \ell \pmod p$. This should explain our choice $\ell = p-1$, $k = 2p-1$, and $\delta = \sqrt{2p}$ for our parameters.

So $\binom{n}{p-1}$ is an upper bound on the independence number of H_n in UD_n^δ and hence

$$\frac{\binom{n}{2p-1}}{\binom{n}{p-1}} = 2^{\left(H\left(\frac{2}{c}\right)-H\left(\frac{1}{c}\right)+o(1)\right)n}$$

is a lower bound on the chromatic number for large $n = cp$. Optimizing the constant factor gives $c = 4 + 2\sqrt{2} \sim 6.82$ and the lower bound is then at least 1.2^n. □

Borsuk's Conjecture

Dead at the age of 60. Died after no apparent signs of illness, unexpectedly, of grave combinatorial causes.

Epitaph of Babai and Frankl for Borsuk's Conjecture

Dividing a task, a structre, a set into smaller, simpler, more transparent parts is *the* number one general approach to solving problems concerning them. In this section we will consider partitioning subsets of the euclidean space into "smaller" ones.

The *diameter* is one of the basic measures of how large a set $B \subseteq \mathbb{R}^d$ is. It is defined as

$$diam(B) := \sup\{||x - y|| : x, y \in B\}.$$

A set is called *bounded* if its diameter is finite.

A long-standing conjecture of Borsuk was concerned about breaking bounded sets into sets of smaller diameter. A partition of the regular d-dimensional simplex $\sigma_d \subseteq \mathbb{R}^d$ into sets of smaller diameter needs at least $d+1$ sets. Indeed, any two of the $d+1$ vertices of σ_d have distance that is equal to the diameter of σ_d. Consequently no two of the vertices can be in the same part of the partition. It is also not difficult to see that a partition of σ_d into $d+1$ sets of smaller diameter is possible.

Borsuk conjectured that the simplex should be a "worst case" with respect to this measure.

Conjecture 0.1.19 (Borsuk's Conjecture (1933))**.** *Every set in $\mathbb{R}^d$ with bounded, non-zero diameter can be partitioned into $d+1$ sets of smaller diameter.*

It is easy to verify the conjecture in dimension 1. Cover a bounded set $B \subseteq \mathbb{R}$ with the two closed intervals $\left[\inf B, \frac{\inf B + \sup B}{2}\right]$ and $\left[\frac{\inf B + \sup B}{2}, \sup B\right]$, each with diameter $\frac{diam(B)}{2} < diam(B)$. [4]

Over the years Borsuk's conjecture was proved for:

- all bodies in dimensions 2 and 3,
- centrally symmetric bodies,
- bodies with smooth surface.

It turned out however, that the general conjecture is not only false, but *very false.*

Theorem 0.1.20 (Kahn-Kalai, 1992)**.** *There exists a bounded set in $\mathbb{R}^d$ that cannot be partitioned into $1.2^{\sqrt{d}}$ sets of smaller diameter.*

Proof. Partitioning a finite point set $P_d \subseteq \mathbb{R}^d$ into sets of smaller diameter means partitioning P_d into independent sets of the δ-distance graph UD_d^δ, where $\delta = diam(P_d)$ is the *largest distance* that occurs among the points of P_d. We will construct a point set $P_d \subseteq \mathbb{R}^d$, such that the largest independent set of UD_d^δ in P_d is small, hence one needs many of them to cover P_d.

[4] The partition is then the intersection of the set B with these intervals.

In our previous theorem, we have already proved a statement in a similar vein for the point set $B_n^{(k)}$ (containing vectors of length n with k 1-entries and $n-k$ 0-entries). For any prime p and uniformity $k = 2p-1$ we have shown that the largest independent set of $UD_n^{\sqrt{2p}}$ contained in $B_n^{(k)}$ is at most $\binom{n}{p-1}$. So the only difference compared to what we want here is the distance we focus on: instead of the *largest* distance, our theorem above focused on the *particular* distance $\sqrt{2p}$.

Our plan is to embed $B_n^{(k)}$ injectively into some $B_d^{(q)}$, so that the pairs of points at distance $\sqrt{2p}$ embed to pairs of points at distance that is the diameter of the image point set P_d. Then

$$\alpha\left(UD_d^{diam(P_d)}\left[P_d\right]\right) = \alpha\left(UD_n^{\sqrt{2p}}\left[B_n^{(k)}\right]\right) \leq \binom{n}{p-1}.$$

How should we define the magic map? In the previous subsection we have already observed how the distance between two characteristic vectors of sets of size q is connected to the size of their intersection: $||\mathbf{v_A} - \mathbf{v_B}||^2 = 2q - 2|A \cap B|$. That is, the diameter appears as the distance between the characteristic vectors of two members A and B of some q-uniform family $\mathcal{F}$ if and only if $|A \cap B|$ is the *smallest possible* intersection size among the members of $\mathcal{F}$.

Distance $\sqrt{2p}$ in $B_n^{(k)}$ corresponds to intersection size $k - \frac{1}{2}\left(\sqrt{2p}\right)^2 = p-1$ in $\binom{[n]}{k}$. Therefore we need an injective map from $\binom{[n]}{k}$ into $\binom{[d]}{q}$ such that pairs of k-sets with intersection size $p-1$ are taken into pairs of q-sets having the smallest possible intersection size that comes up within the image.

The combinatorial idea is to consider the set of $\binom{n}{2}$ pairs of the base set and to each k-subset of $[n]$ associate the set of pairs that are “crossing” between the set and its complement. Intuitively, two such sets of crossing pairs intersect in as few pairs as possible if the sets themselves overlap both each other and each other's complements as little as possible.

Formally, for every $A \in \binom{[n]}{k}$ we define a subset $S_A \in \binom{\binom{[n]}{2}}{k(n-k)}$ of “crossing” pairs as follows

$$S_A := \left\{T \in \binom{[n]}{2} : |T \cap A| = 1\right\}.$$

We count those pairs that are crossing for both sets A and B:

$$\begin{aligned}|S_A \cap S_B| &= |A \cap B||(A \cup B)^c| + |A \setminus B||B \setminus A| = \ell(n-2k+\ell) + (k-\ell)^2\\ &= 2\ell^2 + (n-4k)\ell + k^2,\end{aligned}$$

where $\ell = |A \cap B|$. This quadratic function is minimized at $\ell = \frac{4k-n}{4} = 2p - 1 - \frac{n}{4}$. In order to make sure that for integers it is minimized exactly at $\ell = p-1$, we set $n = 4p$.

Then for the point set

$$P_d := \left\{\mathbf{v_{S_F}} \in \{0,1\}^d : F \in \binom{[n]}{k}\right\}$$

where $d = \binom{n}{2}$, we have that $||\mathbf{v_{S_A}} - \mathbf{v_{S_B}}|| = diam(P_d)$ if and only if $|S_A \cap S_B| = \min\{|S_C \cap S_D| : C, D \in \binom{[n]}{k}\}$ if and only if $|A \cap B| = p-1$ if and only if $||\mathbf{v_A} - \mathbf{v_B}|| = \sqrt{2p}$.

The size of P_d is $\binom{n}{k}$ and hence for large d one needs at least

$$\frac{\binom{4p}{2p-1}}{\binom{4p}{p-1}} = 2^{(1-H(1/4)+o(1))n} = 2^{(1-H(1/4)+o(1))\sqrt{2d}} > 1.1^{\sqrt{d}}$$

subsets of smaller diameter to cover it. □

Bibliography

[1] L. Babai, P. Frankl, Linear Algebra Method in Combinatorics.