6.2 Hypergraphs

In the previous section, we looked at *graphs* that were extremal from some viewpoint. A graph was in fact nothing other than a set of edges, or, in other words, *two-element subsets* of a vertex set.

This concept can be generalized if we allow *larger* subsets. These new collections of subsets are our main topics in the present section.

Definition 6.28 *A* hypergraph *on the set* $[n]$ *is a collection of subsets of* $[n]$*. The elements of this collection, which are subsets of* $[n]$*, are called the* edges *of the hypergraph, while the elements of* $[n]$ *are called the* vertices *of the hypergraph.*

So graphs are just a special case of hypergraphs, namely, the special case in which each edge consists of two vertices. Hypergraphs are often called *set systems* or *families of subsets* as well. They are often denoted by calligraphic letters, such as $\mathcal{F}$.

In this section, we will be looking at hypergraphs that are optimal from some point of view, for instance, they have the highest number of edges possible under certain conditions. For a hypergraph $\mathcal{F}$, let $|\mathcal{F}|$ denote the number of edges of $\mathcal{F}$.

6.2.1 Hypergraphs with pairwise intersecting edges

Theorem 6.29 *Let* $\mathcal{F}$ *be a hypergraph on vertex set* $[n]$ *that does not contain two* disjoint *edges. Then* $|\mathcal{F}| \leq 2^{n-1}$.

Proof: Let us arrange the 2^n subsets of $[n]$ into pairs, matching each subset with its complement. This will result in 2^{n-1} pairs. Each pair consists of *disjoint* edges, so $\mathcal{F}$ can contain at most one edge of each pair. This proves that $\mathcal{F}$ cannot have more elements than the number of these pairs, that is, 2^{n-1}. ♢

The above proof was not very difficult. This might suggest that, with a little extra effort, we can improve the result further. However, that hope is false, as the following theorem shows.

Theorem 6.30 *There exists a hypergraph* $\mathcal{F}$ *on vertex set* $[n]$ *that has* 2^{n-1} *edges so that no two elements of* $\mathcal{F}$ *are disjoint.*

Proof: Let $\mathcal{F}$ be the hypergraph whose edges are the subsets of $[n]$ that contain the vertex 1. ♢

We would like to point out that the hypergraph $\mathcal{F}$ defined in the above proof is actually even better than we needed it to be. It does not only have the property that any two of its elements intersect, but also the property that *all* its elements intersect!

See Exercise 13 for an alternative proof of Theorem 6.30.

One can now ask how large $\mathcal{F}$ can be if we require that any two edges of $[n]$ which are in $\mathcal{F}$ intersect in at least k elements, as opposed to just at least one element. At this point, the reader might suggest that we define $\mathcal{F}$ to be the hypergraph whose edges are the subsets of $[n]$ containing the set $[k]$. That family has 2^{n-k} elements, and it certainly satisfies the criterion on the size of pairwise intersections. The question is whether we can do better. The following construction shows that, if n is large enough, then we actually *can.*

Example 6.31 *Let k be such that $n+k$ is even, and let $\mathcal{F}$ be the hypergraph on vertex set $[n]$ whose edges are the subsets of $[n]$ that have $(n+k)/2$ elements.*

Then any two edges in $\mathcal{F}$ intersect in at least k elements, and if n is large enough, then $|\mathcal{F}| > 2^{n-k}$.

Solution: Indeed, any two edges together have $n+k$ elements, so they must share at least k vertices. On the other hand, the number of edges of $\mathcal{F}$ is $\binom{n}{(n+k)/2}$, and we will show that if n is large enough, then

$$\binom{n}{(n+k)/2} > 2^{n-k}. \tag{6.4}$$

First, note that, if $n = 3k$, then by Stirling's formula (1.13) the left-hand side is equal to

$$\binom{3k}{k} \sim 6.75^k \frac{\sqrt{3}}{2\sqrt{n\pi}},$$

while the right-hand side is just 4^k. Second, note that, if n grows to $n+2$, then the left-hand side grows by a factor of $\frac{2(n+2)(n+1)}{n+k+2}$, whereas the right-hand side simply grows by a factor of four. For large enough n, such as $n = 3k$, the growth rate of the left-hand side is larger, so the left-hand side of (6.4) stays larger than its right-hand side. ◊

Let us now refine our interest and look for optimal hypergraphs in which each edge has the same size. This class of hypergraphs occurs so often that it has its own name.

Definition 6.32 *If all edges of the hypergraph $\mathcal{F}$ consist of k vertices, then $\mathcal{F}$ is called k-uniform.*

First, let us look for large k-uniform hypergraphs that do not contain disjoint edges. Just as for general hypergraphs, we can make sure that the non-disjointness criterion is satisfied by simply placing a given entry into all edges of $\mathcal{F}$.

Proposition 6.33 *For all positive integers $k \in [n]$, there exists a k-uniform hypergraph of $\binom{n-1}{k-1}$ edges on vertex set $[n]$ that does not contain two edges which are pairwise disjoint.*

Proof: The hypergraph of all k-element subsets that contain the vertex 1 has the required property. $\diamondsuit$

It will probably not surprise the reader that the simple construction that we discussed is the best one yet again. While the result is very similar to Theorem 6.29, its proof is considerably more difficult.

Theorem 6.34 **(Erdős–Ko–Rado theorem)** *Let $\mathcal{F}$ be a k-uniform hypergraph on $[n]$ that does not contain two disjoint subsets. Then*

(a) $|\mathcal{F}| \leq \binom{n-1}{k-1}$ *for* $2k \leq n$, *and*

(b) $|\mathcal{F}| \leq \binom{n}{k}$ *for* $2k > n$.

The proof we present is due to Gyula O. H. Katona [48], and it is significantly simpler than the original proof [28] of the theorem.

Proof: Part (b) is true since we can set $\mathcal{F}$ to be the set of *all* k-element subsets of $[n]$. Let us prove part (a).

Think of the elements of $[n]$ as n different people who sit down on n chairs around a circular table that are numbered 0 through $n-1$, as shown in Figure 6.10.

Once a seating is specified, and in that seating k people sit on k consecutive chairs, then we will say that they form a *block* in that seating. More precisely,

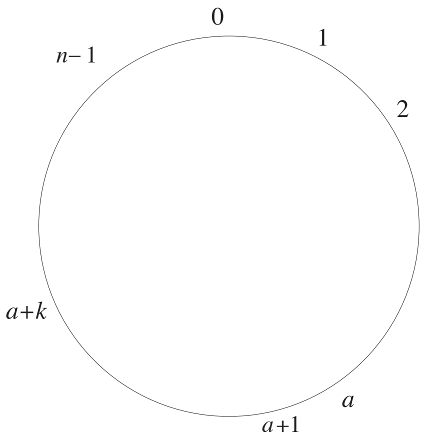

Figure 6.10
Chairs around a table.

if k people sit on chairs $a, a+1, \cdots, a+k-1$ in a given seating, then we will call them the block B_a of that seating. Addition here is meant modulo n, in other words, $n + c = c$.

First, we claim that, for any given seating, $\mathcal{F}$ can contain at most k blocks of that seating. Indeed, let $B_a \in \mathcal{F}$. Then all other blocks in $\mathcal{F}$ have to intersect B_a, so all other blocks in $\mathcal{F}$ have to be of the form B_{a+i} or B_{a-i}, with $i \in [k-1]$. This shows that there are at most $2k-1$ blocks in $\mathcal{F}$, which is not quite what we promised. So we point out that $\mathcal{F}$ cannot even contain all these blocks. In fact, for any $j \in [k-1]$, the blocks B_{a-k+j} and B_{a+j} are disjoint, so $\mathcal{F}$ can contain at most one of them. This shows that, indeed, $\mathcal{F}$ contains at most k blocks.

Now let us count all possible pairs (p, B), where p is a possible seating of our n people around the table and B is a block in that seating that is contained in $\mathcal{F}$. Let P be the number of such pairs.

Let us first count by the seatings. We have just seen that each seating will have at most k blocks which are contained in $\mathcal{F}$. Therefore,

$$P \leq k \cdot n!. \tag{6.5}$$

Let us now count by the blocks. If $B \in \mathcal{F}$, then to seat the people of B in a block of p, first we select the starting chair a of the block in n ways. Then we can seat the k people of B on the chairs $a, a+1, \cdots, a+k-1$ in $k!$ ways, then seat the remaining $n-k$ people on the remaining chairs in $(n-k)!$ ways. Since we can do this for each $B \in \mathcal{F}$, this leads to

$$P = |\mathcal{F}| \cdot n \cdot k! \cdot (n-k)!. \tag{6.6}$$

Comparing (6.5) and (6.6), we get

$$\begin{aligned} |\mathcal{F}| \cdot n \cdot k! \cdot (n-k)! &\leq k \cdot n! \\ |\mathcal{F}| &\leq \frac{k}{n} \cdot \binom{n}{k} \\ &\leq \binom{n-1}{k-1}. \end{aligned}$$

◊

If you think you understand this proof, you can test your understanding by finding the point in the proof at which we used the condition $2k \leq n$.

We would like to point out an important feature of the proof that will be useful many times in the future. The key element of the proof was to count *pairs* (p, B), once by the permutations, once by the blocks. This technique, counting specific pairs, once by their first elements, once by their second elements, and then comparing the results, is a simple but powerful tool.

The Erdős–Ko–Rado theorem has many variations. The following one, in a more general form, was proved by Béla Bollobás [11].

Theorem 6.35 *Let $\mathcal{F}$ be a k-uniform hypergraph on $[n]$ with $2m$ edges so that its edges can all be listed as $X_1, X_2, \cdots, X_m$ and $Y_1, Y_2, \cdots, Y_m$ in a way that $X_i \cap Y_j = \emptyset$ if and only if $i = j$. Then*

$$|\mathcal{F}| \leq \binom{2k}{k}.$$

One could think of this problem as follows: A football coach has n players and wants them to form various teams to play m short games against each other. In each of these games, one team will play offense only, and the other team will play defense only. No player can be on both teams in the same game (hence the $X_i \cap Y_i = \emptyset$ restriction). Furthermore, since this is a combinatorially inclined coach, no offense can be totally disjoint from any defense other than its current opponent, and vice versa, no defense can be totally disjoint from any offense other than its current opponent. (It is not clear that many coaches would worry about violating this last condition, but this one does.)

Note that the upper bound does not even depend on n. No matter how large n is, the size of the best possible $\mathcal{F}$ will not grow. This shows that we are dealing with a much stronger restriction here than in the Erdős–Ko–Rado theorem.

Proof: (of Theorem 6.35) Let us count the number P of pairs (p, i) so that $p = p_1 p_2 \cdots p_n$ is an n-permutation, $i \in [m]$, and all elements of X_i precede all elements of Y_i in p.

First, let us count by the indices $i \in [m]$. Let i be fixed, then X_i and Y_i are also fixed. Let us say that two permutations q and q' are equivalent if the set of entries that belong to $X_i \cup Y_i$ is located in the same positions in the two permutations. Note that we do *not* require that each entry $z \in X_i \cup Y_i$ be located in the same position in q as it is located in q'. What we do require is that the *set* of positions of the entries that belong to $X_i \cup Y_i$ is the same in q as in q'.

Then there are $\binom{n}{2k}$ equivalence classes, and each equivalence class consists of $k! \cdot k!(n-2k)!$ permutations. Indeed, there are $\binom{n}{2k}$ for the set of positions of the entries that belong to $X_i \cup Y_i$. Once that set is chosen, the k elements of X_i have to be in the first half of the positions of that set, in any order, and the k elements of Y_i have to be in the second half, again in any order. Finally, the entries that do not belong to $X_i \cup Y_i$ can be permuted in any order. This argument works for any fixed i, showing that

$$P = m \cdot \binom{n}{2k} \cdot k! \cdot k!(n-2k)! = n! \cdot \frac{m}{\binom{2k}{k}}. \tag{6.7}$$

Now let us count the pairs (p, i) with the required property by the permutation p. Let p be fixed. We claim that there is *at most one* index i so that (p, i) has the required property. Assume not; that is, assume that both i and j are such indices. Assume without loss of generality that the rightmost element of X_i in p is weakly on the right of the rightmost element of X_j in

Figure 6.11
If X_j ends before X_i, then X_j is disjoint from Y_i.

p. That would mean that $X_j \cup Y_i = \emptyset$, which is a contradiction. Indeed, all elements of X_i precede all elements of Y_i in p, so all elements of X_j must precede all elements of Y_i in p, since X_j ends before X_i. See Figure 6.11 for an illustration.

Since there is at most one index i for each permutation p, we certainly have $P \leq n!$. Comparing this inequality with (6.7), we get $n! \cdot \frac{m}{\binom{2k}{k}} \leq n!$, or $m \leq \binom{2k}{k}$ as claimed. ♢

The reader is asked to prove that this upper bound cannot be further improved (Supplementary Exercise 13). And then the reader is asked to generalize the statement of this theorem for nonuniform hypergraphs (Supplementary Exercise 14).

6.2.1.1 Sunflowers

In the constructions that proved to be optimal in the situations discussed in Theorem 6.34 and Theorem 6.30, we saw examples of hypergraphs in which the intersection is nonempty. This concept can be strengthened by requiring that the intersection of any two edges be not only nonempty, but also the *same*. This class of hypergraphs is so important that it has its own name.

Definition 6.36 *A hypergraph $\mathcal{H}$ of nonempty edges is called a* sunflower *if each pair of edges of $\mathcal{H}$ has the same intersection, which is called the* core *of the sunflower. The edges of the sunflower are called the* petals*.*

Example 6.37 *The hypergraph consisting of edges $\{1, k\}$, where k ranges over the set $\{2, 3, \cdots, n\}$ is a sunflower. See Figure 6.12 for an illustration of this sunflower for $n = 9$, and an explanation for the name "sunflower."*

Note that the core is allowed to be empty. Also note that Definition 6.36 is equivalent to saying that the intersection of any two edges is equal to the intersection of *all edges*.

For instance, if the edges of $\mathcal{F}$ are pairwise disjoint, then $\mathcal{F}$ is a sunflower, even if that is not a particularly exciting example. What is more interesting is that large-enough hypergraphs always contain sunflowers. This is the content of the famous Erdős–Rado lemma. Let us say that $\mathcal{H}$ is a (p, ℓ)-*sunflower* if $\mathcal{H}$ is a sunflower with p petals, none of which consists of more than ℓ vertices.