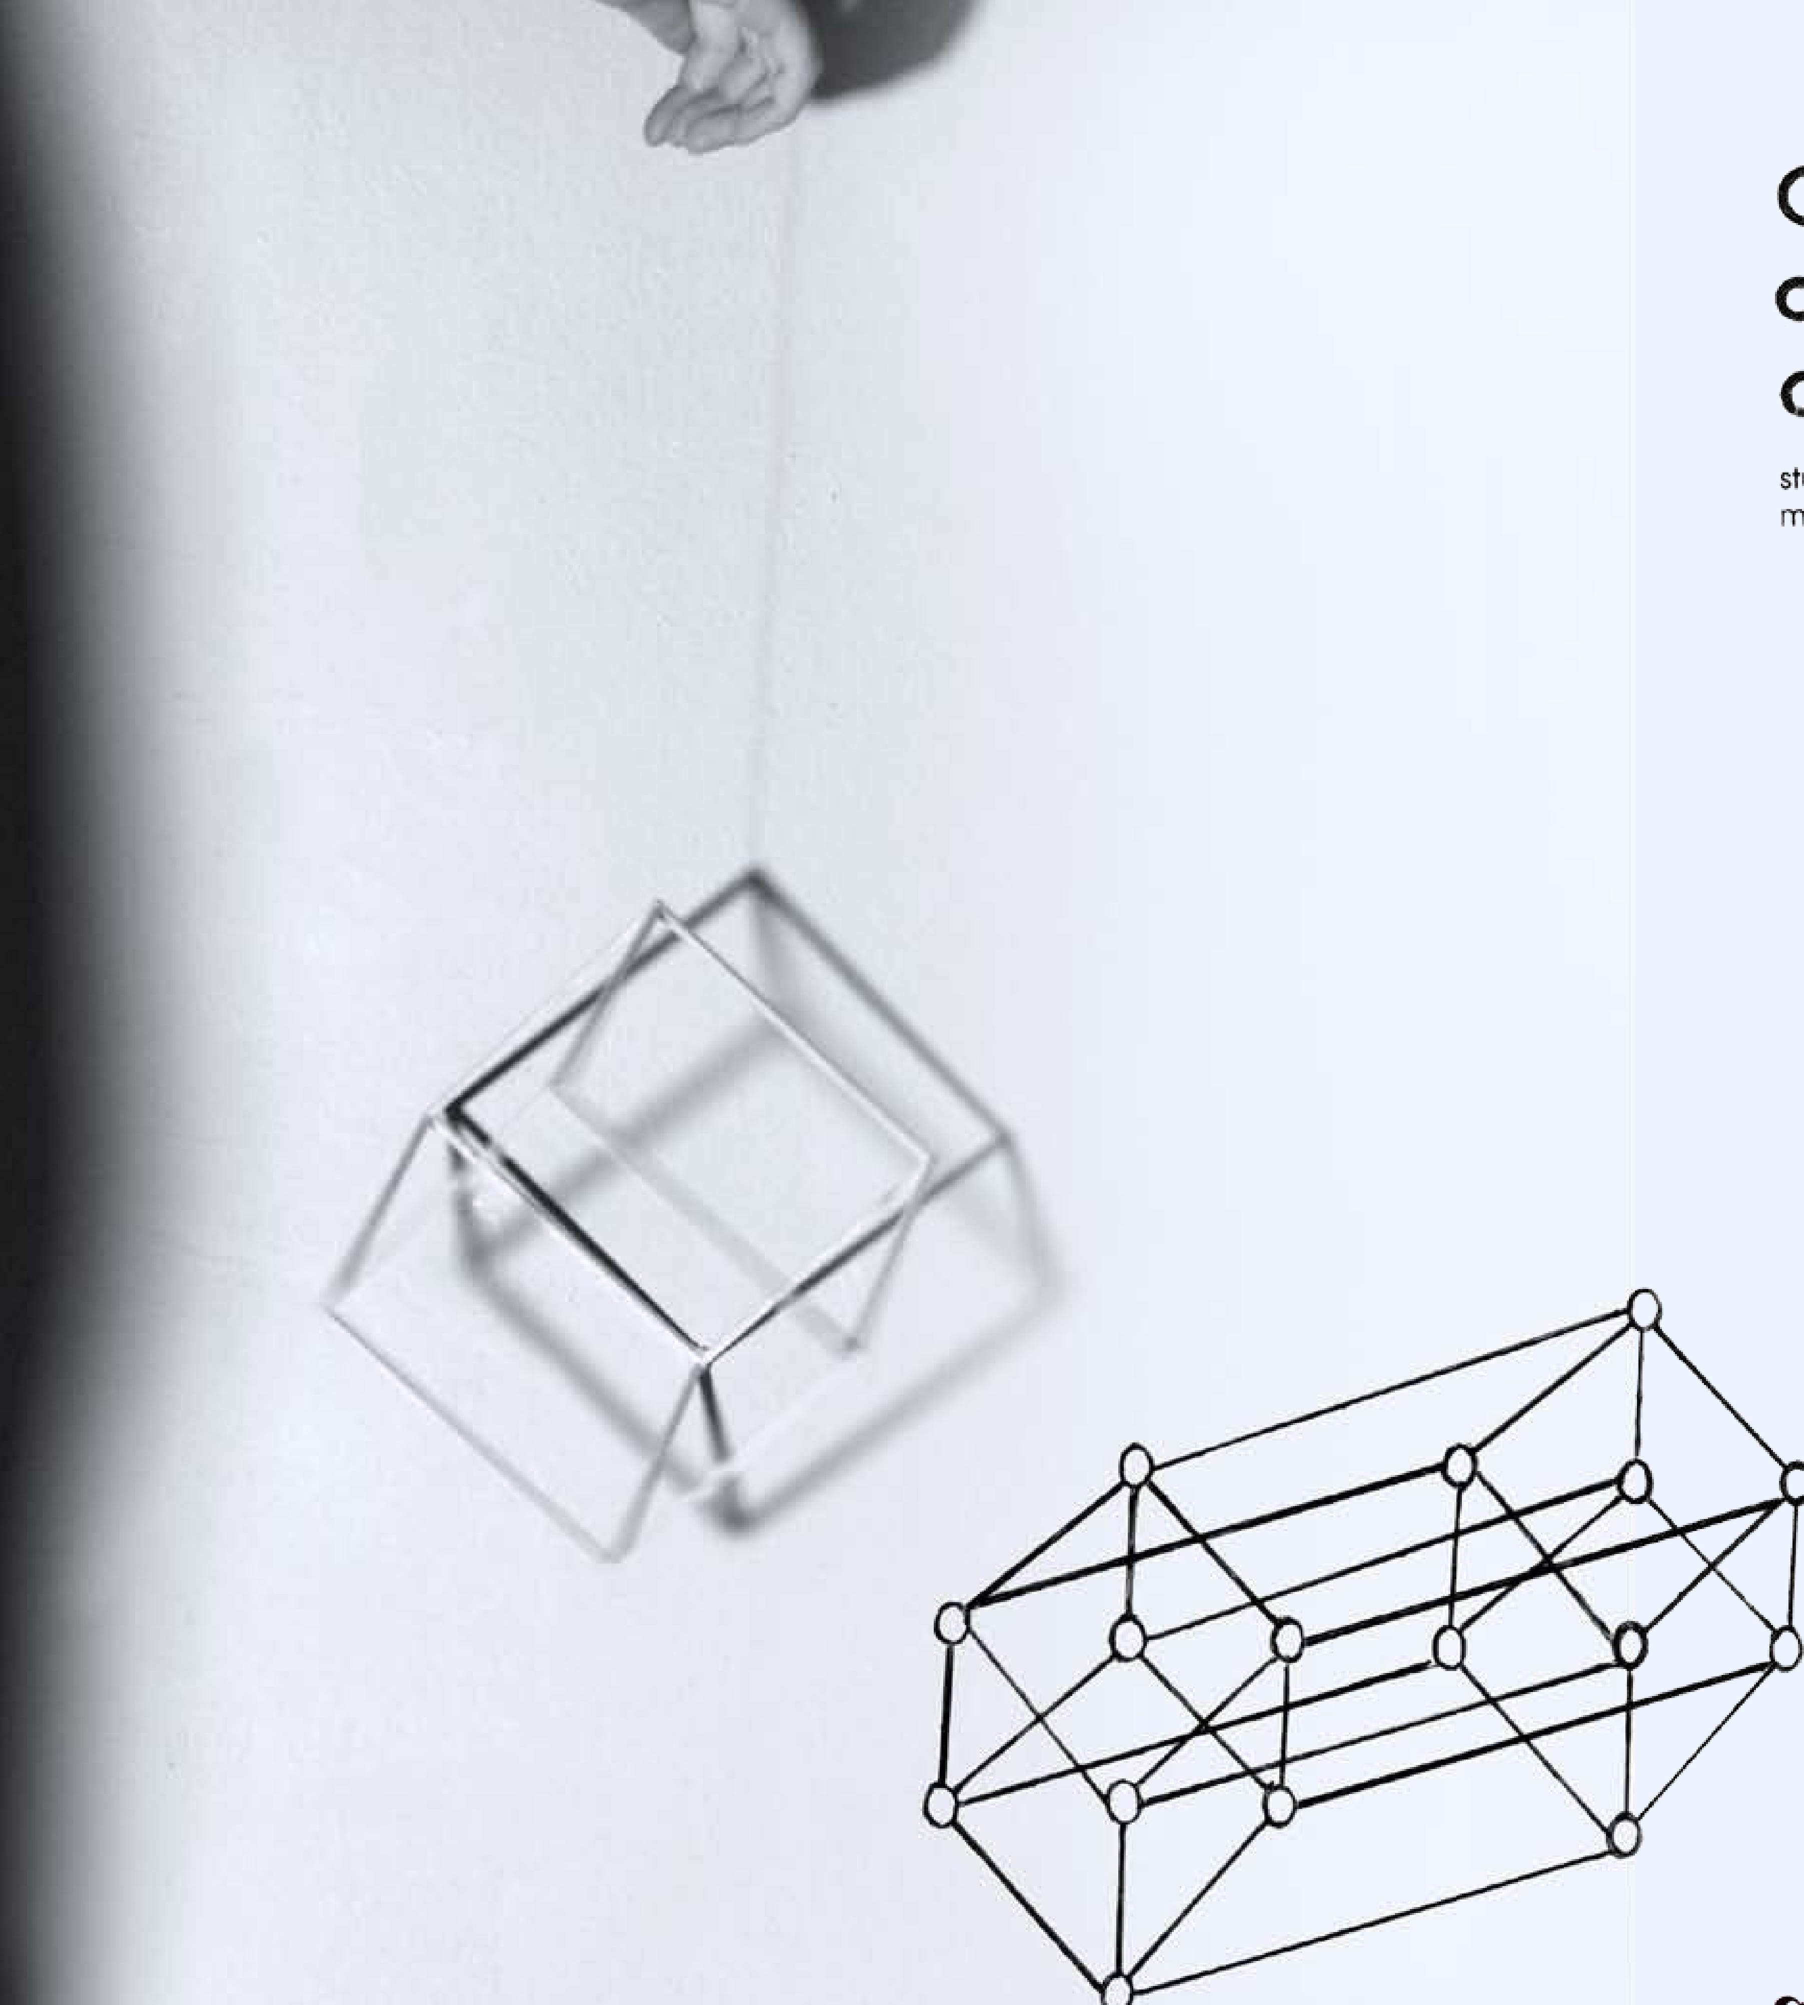

Combinatorial Optimization over Maximal Intersecting Families of Subsets

student: David Renshaw
mentor: Cheng-Yeaw Ku

One of the most basic structures in mathematics is that of a finite set and its subsets. Proving theorems without it is like baking cakes without a mixing bowl—impossible because there is nowhere to put anything. Therefore it is no surprise that combinatorialists, who specialize in pointed questions about elementary structures, have often studied collections (affectionately called `families') of subsets satisfying various criteria. One particularly natural question, however, has remained unsolved for more than twenty years. Roughly, it asks whether a certain optimization problem concerning families of intersecting subsets (of a given finite set *S*) always has a simple solution. An affirmative answer would imply the resolution of another older, open question that had stumped even the legendary mathematician Paul Erdös. The present paper documents an attempt to use one specific idea—what we will call `reductive smoothing'—to attack this problem. Although the method apparently leads to a dead end, it yields some insights along the way that are, perhaps, interesting in their own right.

An Intuitive Conjecture

The problem itself is easily stated. The only notion that requires special definition is that of a `maximal intersecting family' of subsets of *S*. Here, `intersecting' means that any two members of such a family share at least one common element (not necessarily the same for all pairs) and `maximal' means that if the family adopts any new members it will cease to be intersecting. One important fact about these families is that each consists of exactly half of the subsets of *S*. Now suppose that every subset of *S* is assigned a weight, with the restriction that each subset must not weigh more than any of its subsets. That is, suppose weight is non-increasing with respect to inclusion. (The older question alluded to above assumes furthermore that weight can only take the values *0* or *1*.) Then suppose that we would like to find a maximal intersecting family whose total weight is as large as possible. In 1979, Daniel Kleitman conjectured that the only families we need to check are those which contain a singleton—that is, a set of one element. The aim of our research was to prove this conjecture.

Intuitively, the validity of the conjecture might seem almost maddeningly obvious. Families that contain a singleton are in a sense `pushed down' as far as they will go (the figures provide an explanation of this terminology). Since smaller *subsets* are heavier, why should not the heaviest *family* contain as many of them as possible? To put it another way, imagine a maximal intersecting family that does not contain a singleton but has maximum weight. There will be some subset *A* that is a member but that contains no other member. We can obtain another maximal intersecting family by trading *A* for its complement A^c, so the weight of *A* must be greater than the weight of A^c. But then should we not be able to make a bunch of favorable trades to include all of those heavy subsets contained in *A*? Or if heavy members somewhere else prevent us from doing this, should we not be able to make favorable trades to include *their* subsets? In either case, we will ultimately be led to a family that contains a singleton. Making this line of reasoning rigorous, however, is not easy.

A Smooth Approach

Fortunately, we do not need to consider every possible weight function individually. A technique introduced by Kleitman and later elaborated upon by Peter Fishburn allows us instead to focus our attention on the maximal intersecting families themselves. The idea, which we will call `smoothing,' is to take an arbitrary family and show that it can never be lighter than some average of singleton-containing families. When smoothing succeeds, it demonstrates that one of the singleton families must weigh at least as much as the original family. Hence, if we can prove that every maximal intersecting family can be smoothed, then we will have proved the conjecture. In fact, if we ever find a family that cannot be smoothed then we will actually have disproved the conjecture.

This equivalence (which we will not prove here) is encouraging.

The smoothing process can be illustrated by an example. Suppose that F is a maximal intersecting family which contains the doubleton $\{a,b\}$. Each member of F will then fall into one of three disjoint classes, depending on whether it contains a, b, or both. Let s_a, s_b, and s_{ab} be the respective sizes of these classes. We know right away that $s_{ab} = 2^{n-2}$ because that is how many supersets $\{a,b\}$ has, and all of them must be in F by maximality. This leaves another 2^{n-2} members of F that are split in some way between the other classes. It turns out that the total weight of F must be less than or equal to $s_a/2^{n-2}$ times the weight of the family containing the singleton $\{a\}$, plus $s_b/2^{n-2}$ times the weight of the family containing $\{b\}$. To see why this is so, consider the following interpretation. Imagine that we have written down all of the subsets of our ground set, and we place a one-ounce ball of clay on each subset that is a member of F. To calculate the `weight' of this configuration we multiply the real weight of each ball of clay by the value of the weight function on the subset it covers. Now, whenever we transfer clay from some subset to a subset contained in it, the weight of the configuration will remain constant or increase. In fact, it is not too difficult to show that we can always transfer clay in this favorable way so as to evenly distribute the s_a class-a ounces of clay among *all* subsets that contain a but not b, and the s_b class-b ounces among all subsets that contain b but not a. Furthermore, each of the s_{ab} balls still intact after this can be split into a $s_a/2^{n-2}$ ounce portion and a $s_b/2^{n-2}$ ounce portion. It is now clear that the weight of F is no greater than the weight of an average of singleton families.

Our idea is to generalize this technique. Evidently, smoothing succeeds when F has a doubleton because then we can forget about all but two singleton families. It would be helpful in the general case if we could forget even about just one singleton family. There turns out to be a nice property which, when exhibited by a family, guarantees that we can indeed do this. For any maximal intersecting family F and any element x of the ground set S, we may consider the induced family $F^{(x)}$ obtained by deleting every instance of x from the members of F. Of course, $F^{(x)}$ is not necessarily a maximal intersecting family, and in general it is not even intersecting. However, if the members of $F^{(x)}$ can be partitioned into two classes, each of which is intersecting, then the problem of smoothing F reduces to the problem of smoothing two families on the smaller ground set $S\backslash\{x\}$. In other words, by forgetting about the singleton family containing $\{x\}$, we can reduce to a smaller problem and hence allow for an inductive proof of the conjecture.

Our Optimistic Hypothesis

Which maximal intersecting families allow such `reductive smoothing?' Clearly, if they all do, then a proof of the conjecture follows immediately. Unfortunately, the Fano plane dashes any such hopes. It is a family of three-element subsets of a seven-element set such that any two members intersect in exactly one element and each element is contained in exactly three members. It fails the test because any family induced from it by deleting an element necessarily contains three disjoint members. Nevertheless, the Fano plane poses us little difficulty because the unique maximal intersecting family generated by it can be *evenly smoothed*; if we place a ball of clay on each member of the Fano plane and on each superset of these members, we can then favorably transfer clay to obtain the configuration where each subset is covered by an amount of clay proportional to its size.

We conjectured that this will always be so: for every maximal intersecting family, either reductive smoothing or even smoothing will succeed. The observation here is that a family F which cannot be reductively smoothed is forced to have a certain degree of regularity. Namely, for each x in S, there must be a cycle of odd length of members of F where adjacent members intersect only in x. Contrapositively, families that cannot be evenly smoothed seem unlikely to have such a cycle. Consider the case (which will be our focus for the rest of the paper) when S has an even number n of elements and F has no member with less than $n/2$ elements. In this situation F is `pushed up' as far as possible, and thus can *probably* be smoothed evenly. One exception is when there is an element x that is not in any size-$n/2$ member of F. Then there are exactly as many non-members containing x as members containing x. But the x-containing members are the only allowable sources of clay for these non-members, and it turns out they cannot provide quite enough for an even smoothing. However, this family F is precisely the easiest family to reductively smooth, because $F^{(x)}$ is itself intersecting. So perhaps the two kinds of smoothing take care of every family in this situation.

How Necklaces Throw in a Wrench

Surprisingly, our wishful thinking meets with a counterexample even in this simplified setting. Since a `pushed-up' F contains no subsets of size less than $n/2$, all subsets of size greater than $n/2$ are members. This means that our only freedom lies in choosing which size-$n/2$ subsets to include in F. The key to our following construction is that we can choose a small number of these to prevent reductive smoothing, and we can choose the rest so as not to contain some element x, which will still prevent even smoothing. Imagine that we have a string and n beads labeled as elements of S. We choose one bead to set aside and we string up the rest and tie a loop to make a necklace. There are $n-1$ ways to choose $n/2 - 1$ alternating beads on the necklace. Let each of these ways, along with unstrung bead x, define a member of F. This guarantees that $F^{(x)}$ cannot be partitioned into two intersecting families. We now untie the necklace and proceed to set aside a different bead. This time we have to be careful when stringing up the necklace because some orderings will define subsets that are disjoint from those we have already added as members of F. But it turns out that if n is large enough then we can always find a good way to do it, and we

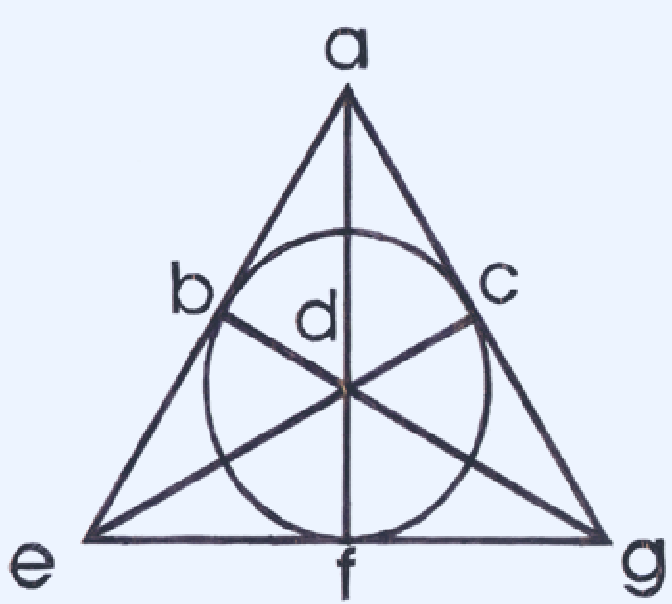

The Fano Plane: {{a,b,e}, {a,d,f}, {a,c,g}, {b,d,g}, {c,d,e}, {b,c,f}, {e,f,g}}

An illustration of the smoothing process.

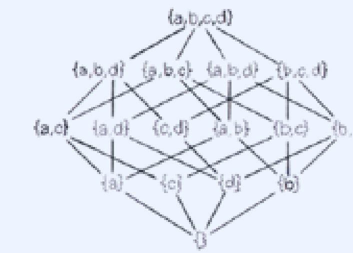

First we write down all subsets of the ground set. We draw a line segment connecting subsets which differ by only one element.

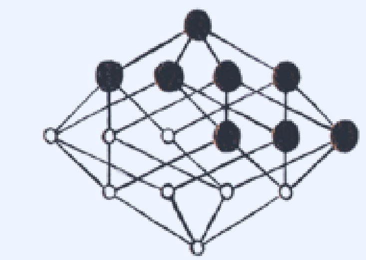

Now we place a ball of clay on each subset which is in our fmaily.

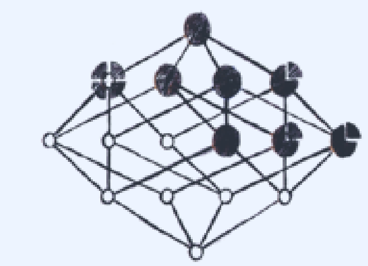

We divide some of the balls intp pieces...

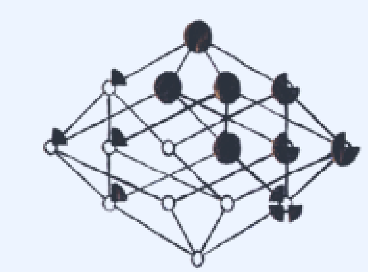

...and move some clay down along line segments. The weight of the configuration can only increase during this step.

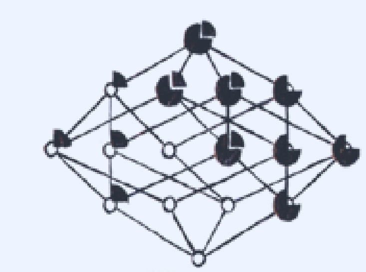

Finally we notice that the resulting configuration is an average of singleton families—in this case. 1/4 of the 'a' family and 3/4 of the 'b' family.

An explanation of 'pushed up' and 'pushed down.'

subsets of size:

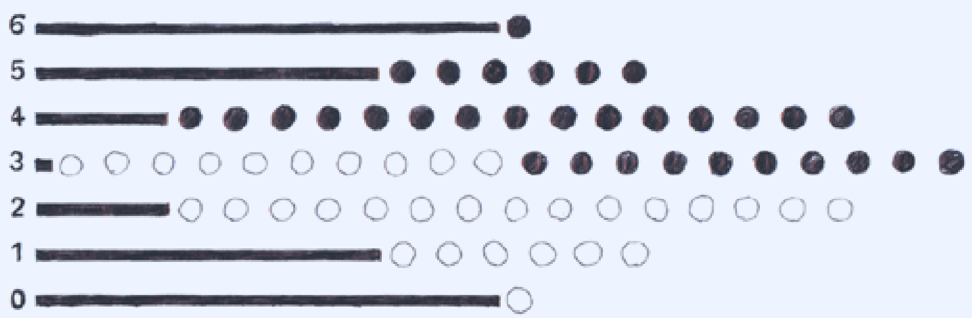

A 'pushed up' family for n=6. We hypothesized that any such family would either be evenly or reductive smoothable.

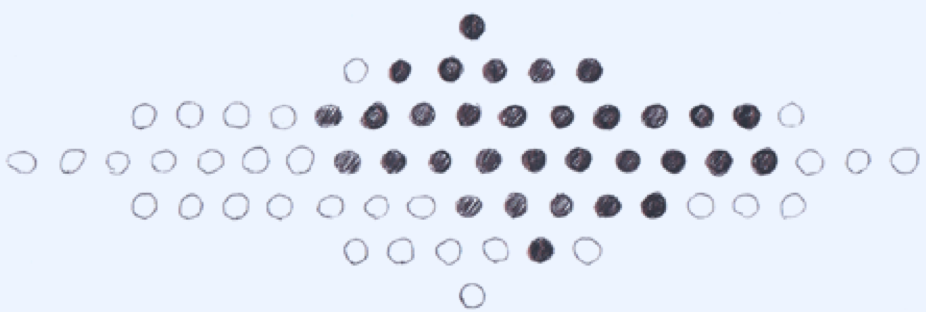

A singleton family for n=6. Such a family is 'pushed down' as far as possible.

can in fact continue until we have set aside each bead. The reason is that although each subset we add to F eliminates a number of feasible necklaces, the total number of ways to make a necklace out of $n-1$ beads is large enough to make this irrelevant. When we are finished, we will have chosen $n(n-1)$ or less members of F that together prevent a reductive smoothing. As mentioned before, we choose the rest of the members of F so as not to contain x. Now, if F contained *no* size-$n/2$ subsets containing x and we tried to evenly smooth, the total clay deficit for the nonmembers containing x would be (as it turns out) exactly $2^{n-2}/n$ ounces. Each size-$n/2$ member of F that *does* contain x can only help this deficit by one ounce. As long as we have chosen n large enough so that $2^{n-2}/n - (n-1)n$ is positive, we find that F cannot be evenly smoothed. Since it also cannot be reductively smoothed, F is our counterexample.

Thus, our approach evidently fails to provide a breakthrough. Reductive smoothing is not strong enough to take care of all families that cannot be evenly smoothed. Does this mean that the technique is worthless? Probably not. Perhaps there is some nice property other than "evenly smoothable" that can be used to fill the gap left by reductive smoothing. Or perhaps we did not define reductive smoothing itself in the most general way. In fact, there do exist *non*-intersecting families of size 2^{n-1} that are nonetheless smoothable. This could possibly be taken into account somehow. In any case, the study of reductive smoothing isolates some of the difficulties inherent in the effort to prove Kleitman's conjecture. And that is the first step towards overcoming them.

Acknowledgements

I am thankful for generous support I received from the Richter Memorial Funds. I would also like to thank my mentor, who provided a sense of direction to my research, and my family, who have been a constant source of encouragement.

Further Reading

1. V. Chvátal. "Intersecting families of edges in hypergraphs having the hereditary property." Hypergraph Seminar. C. Berge and D. Ray-Chaudhuri, eds. Lecture Notes in Math. 411, Springer-Verlag, New York, 1974, pp. 61-66.
2. P. C. Fishburn. "Combinatorial optimization problems for systems of subsets." SIAM Review 30: 4 (1988), 578-588.
3. D. J. Kleitman. " Extremal Hypergraph Problems." Surveys in Combinatorics, B. Bollobas Ed., Cambridge University Press, Cambridge, 1979, pp. 44-65.