Convex Sets in the Plane with Three of Every Four Meeting

Daniel J. Kleitman
Massachusetts Institute of Technology
e-mail: djk@math.mit.edu

*András Gyárfás**
Computer and Automation Research Institute of the Hungarian Academy of Sciences
e-mail: gyarfas@luna.aszi.sztaki.hu

Géza Tóth†
Massachusetts Institute of Technology and Hungarian Academy of Sciences
e-mail: toth@math.mit.edu

Abstract

Suppose we have a finite collection of closed convex sets in the plane, (which without loss of generality we can take to be polygons). Suppose further that among any four of them, some three have non-empty intersection. We show that 13 points are sufficient to meet every one of the convex sets.

1 Introduction

We consider the following simple problem. Suppose we have a finite collection of closed convex sets in the plane, (which without loss of generality we can take to be polygons). Suppose further that among any four of them, some three have a point in common. How many points are required to meet every member of the collection? This is the simplest non-trivial case of a general problem: suppose in d dimensions we have convex polyhedra with the property that among any k of them there are j (for j at least $d+1$) with a point in common. How many points are necessary to meet every member of the collection? That there is a finite answer, independent of the size of the collection was a conjecture of Hadwiger and Debrunner [HD57]. That there is a finite upper bound, $f(k,j,d)$, for the answer to this problem was shown by Alon and the first author, a few years ago [AK92], [AK97]. The argument then given was general but not very efficient, in the sense that the best bounds obtainable from it are rather high. For the first question above using simple tricks to improve it one can reduce it to approximately 200. On the other hand the correct

*Supported by grant OTKA-T-029074.
†Supported by NSF grant DMS-99-70071, OTKA-T-020914 and OTKA-F-22234.

answer to this particular problem may well be 3. This argument had the additional peculiarity that it was based on a counting argument in which the geometric features of the problem appeared only through general concepts; namely the fractional version of Helly's theorem [KL79], Wegner's Theorem [W75] and Epsilon-Net constructions [ABFK92].

The purpose of the present paper is to provide a direct argument for the simple planar problem described in the first paragraph above. This argument has the advantage of giving a much better upper bound, namely 13, which can probably be improved with more ingenuity. We have made no attempt to generalize this result at all, but express the hope that special as it may seem, it has aspects that may be generalized in some direction. It is our belief that it can be refined and improved, but have found the problem so elusive that we think the present result, which is in the end quite simple and elementary, is worth presenting as it stands. We approach the problem by developing a series of simpler problems, of growing difficulty, culminating in the original problem. We will show that each can be used as a tool to solve the next.

2 Convex sets each of whose pairs have a meeting point on one of two lines

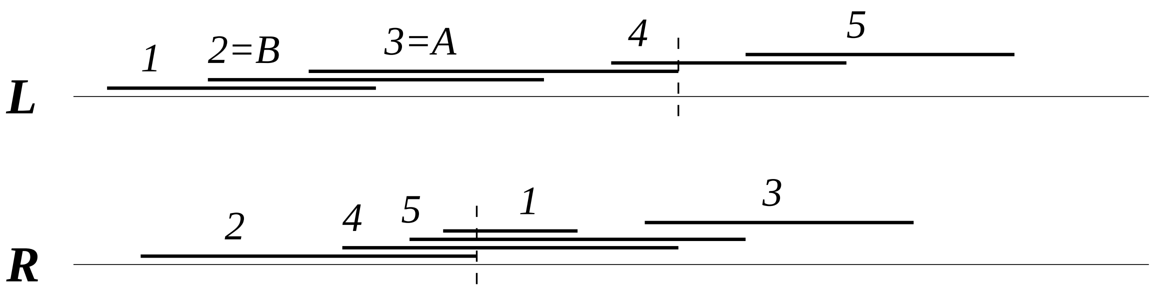

Figure 1. Objects represented by intervals on two lines.

We begin with the following simple observation. Both of the results of this section are special cases of similar problems on trees that have been discussed by A. Gyárfás and J. Lehel [GL70] (see also [GL85]).

Observation 1: Suppose we have two line segments, L and R, and a collection of objects $\{A_i\}$, such that each A_i is represented by a closed interval (denoted by A_i^L) on L and by the closed interval A_i^R on R. Suppose further that for every i and j, either A_i^L and A_j^L have a point in common or A_i^R and A_j^R have a point in common, or both pairs do so. Then two points, one on L and one on R, suffice to meet at least one representor for each object.

Proof: We order L and R linearly from one end of each to the other. Let A be the object whose L-representor ends first on L, among those objects whose representor on R is disjoint from that of another object on R whose representor ends before it on L. Let B be the object whose representor is furthest from A on R among those whose representor on L ends before that of A on L. Then the right endpoint of A^L and the endpoint of B^R closer to A^R meet every object (see Fig. 1).

This complicated sounding paragraph can be expressed simply if we represent our objects by rectangles each of whose horizontal component is its representor on L, and its vertical component its representor on R. (If there is any object whose representor on either L or R is empty, we can give it a fictitious interval that meets no other interval without changing any results; and then the object will be represented by a rectangle here.) In rectangular terms, our condition is that every pair of rectangles have a common vertical or horizontal line. We will show that there is a horizontal and vertical line which between them meet every object.

Consider all pairs of our rectangles that have no common horizontal line, and pick out the pairs that end (on moving rightward) at the leftmost line, call it rA, among such pairs. Let A be the rectangle that ends at rA within such pairs. Among those rectangles that are partners of A in such pairs, let B be the one vertically furthest from A, and let its horizontal end line nearer to A be nB.

From now on, we say that an object meets a line if its represented rectangle does. We divide the objects into three classes. Those that meet rA, those whose right end is to the left of rA, and those whose left end is to the right of rA. Those whose left end is to the right of rA share no vertical line with either A or B and therefore must share horizontal lines with both A and B, so that they must meet nB. Those whose right end is to the left of rA must have a common horizontal line with B by the definition of A, and must extend closer to A vertically so by the definition of B, again meet nB. If there is no such A, then all objects meet in R and a single point in R suffices to meet them all (see Fig. 2).

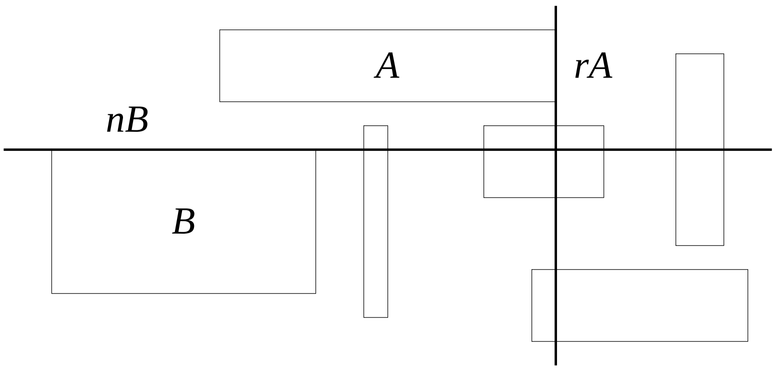

Figure 2. Rectangles with each pair having a common horizontal or vertical line.

3 The Special Configuration

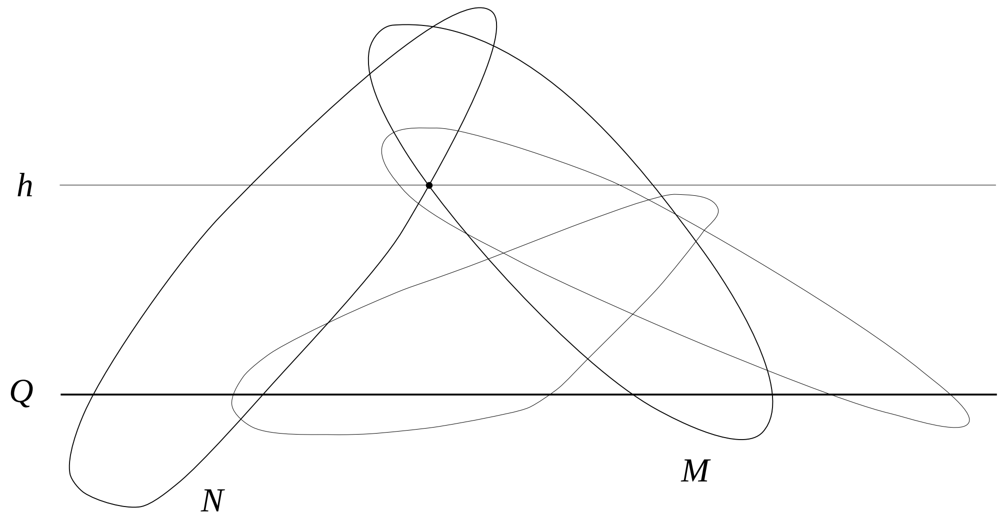

Figure 3. A Special Configuration.

We now consider collections of convex polygons which, while still special, bear a much closer relation to our original problem than the problem considered so far do. We will be able to meet all the polygons in them by judicious use of the previous result. We first define these configurations, which we will call special configurations of polygons. Suppose we have a horizontal line, Q. Then a special configuration of convex polygons is one with the following properties.

1. Each pair of polygons meet on or above Q; each polygon meets Q.
2. Any three polygons that are mutually disjoint on Q have a point in common above Q.

Obviously we can replace above Q by below without any change in the result here:

Theorem 1: *All polygons in a special configuration can be met by most* 5 *points.*

Proof:

Let $\mathcal{A}$ be a a special configuration of convex sets. For a set S of points, $\mathcal{A}(S)$ is the subfamily of $\mathcal{A}$ not pierced by the points of S. For $N, M \in \mathcal{A}$, I_{NM} is the convex hull of $N \cap Q$ and $M \cap Q$. Line h is the line parallel to Q at height h.

We assume that our polygons are jiggled, so that at any given point at which two polygons separate, no other pair of polygons also separates unless one polygon comes to an end. Such jiggling can be performed without changing any polygon intersection properties.

Two polygons, N and M will be said to separate at height h, if they have a point in common at height h above Q and no such point for any smaller height (see Fig. 3). After jiggling this will happen one pair at a time. It can happen that the intersections of many pairs of polygons end at the same point, if one polygon ends there. Since each of our polygons meets Q, this cannot happen above Q here.

We order pairs of polygons in decreasing order by their separation heights, and for each pair of polygons (N, M) in this order, we consider the pair of intervals on Q defined by $N \cap Q$ and $M \cap Q$ as representors of the object (N, M) on Q. These intervals are necessarily disjoint on Q for positive h.

Let the maximal initial segment be defined by the pairs (N, M) for which the intervals I_{NM} pairwise intersect. Let $I = [p_1, p_2] \subseteq Q$ denote the interval of their intersection. Clearly I is the intersection of two *support intervals* $I_{N_1' M_1'}$ with endpoint p_1 and $I_{N_2' M_2'}$ with endpoint p_2 where (N_i', M_i') are (not neccesarily distinct) pairs from the initial segment.

Assume first that the maximal initial segment contains all pairs separating at some positive height. Then $A \in \mathcal{A}(p_1)$ can be partitioned into two parts, those which intersect Q in $(-\infty, p_1)$ and those which intersect Q in (p_1, ∞). Both parts intersect Q in pairwise intersecting intervals therefore three points of Q pierce $\mathcal{A}$.

Otherwise, the maximality of the segment is witnessed by two pairs: (N'', M'') separating at line h'' next to the initial segment and the pair (N_i', M_i'), separating at height h_i' whose support interval is separated from $I_{N''M''}$ by the endpoint p_i of I ($i = 1$ or $i = 2$). Without loss of generality, assume that $i = 1$. The intervals $I_{N''M''}$ and $I_{N_1' M_1'}$ are disjoint on Q.

Now the sets of $\mathcal{A}(p_1)$ can be partitioned into two parts, $\mathcal{A}^-(p_1)$ and $\mathcal{A}^+(p_1)$, that is, those which intersect Q within $X_1 = (-\infty, p_1)$ and those which intersect Q within $X_2 = (p_1, \infty)$, respectively. Let C and D be two sets in $\mathcal{A}^-(p_1)$ that are disjoint on Q, and intersect Q in $X_1 = (-\infty, p_1)$. Then C, D, N_1' and M_1' are all disjoint on Q, so by our initial condition and Helly's theorem, they must have a common point, which must necessarily lie above h''. But then C and D must meet on the line h'', since they did not separate above h''. We conclude that any pair, (C, D), of $\mathcal{A}^-(p_1)$ either meet on Q or on the line at height h'' parallel to Q. By Observation 1, these can all be met by two points, one on Q, one on h''. We can proceed similarly for the sets in $\mathcal{A}^+(p_1)$, the only difference is that we use N'' and M'' in the argument, instead of N_1' and M_1'.

Thus, with p_1, five points pierce $\mathcal{A}$.

4 Reduction to Special Configurations

In this section we show that given a collection of closed and finite convex sets, with three out of any four having a point in common, we can choose three points, p, p' and p'', so that those of our sets that contain none of them form at most two special configurations, as described in the previous section. This implies that we can cover all the sets with 13 points, five for each special configuration and the first three. The initially given (any three of four) condition allows two cases: either all pairs of sets have non-empty intersections, or not. In the second case, no two disjoint pairs of sets, (A, B) and (C, D) can each have empty intersection, or else no three of (A, B, C, D) could have a common point. This means (apart from the degenerate case of three pairwise disjoint sets) that the empty intersection pairs form a "claw".

Suppose sets A and B have a point in common, but their intersection has no point in common with any third of our sets. Then as far as our three out of four condition is concerned, A and B may as well be disjoint, and we can conclude from it that every other pair (C, D) must together have a common point with either A or B or both. In this circumstance we say that A and B are effectively disjoint.

We choose our first point so that we can find two sets A and B that are effectively disjoint, and all other sets meet both of these.

In the first of the two cases above, when all sets meet, we can choose an arbitrary direction, $\bar{n}$, and move a line normal to $\bar{n}$ from infinity until we encounter the first point p, at which a pair of our sets, A and B, separate. p will be our first point, and as far as the other sets that do not contain p are concerned, A and B are effectively disjoint.

In the second case, when there is let the hub of the claw that represents the disjointness graph among our sets be the set A, and the sets that it does not meet be $B_1, B_2, \ldots$. Then any triple among the B_i's have a point in common as we can deduce by applying the three of four condition to the triple and A. Then Helly's Theorem in the plane implies that all the B_i's have a point in common. If we choose any common point as p, then A and any one of the B_i's, form two sets sought here; they are actually disjoint, and all sets not containing p will meet both.

In the first case, let Q be a line through p which touches A. In the second, let Q be a line that separates A and B and touches A. We will refer to Q as a horizontal line.

We choose our second and third points, p' and p'', so that the remaining sets not containing either of them can be partitioned into two special configurations, one in which all pairs of members meet above Q, while those in the other all meet below Q. We define these points explicitely at the end of Section 7.

By considering all four-tuples of the form (A, B, X, Y) among our sets X and Y, because A and B are effectively disjoint we can deduce from our three of four condition that every pair, X and Y, must meet either in A or in B.

We consider the complete graph on our convex sets as vertices and label each edge of that graph with one of three labels: (X, Y) gets label A if X and Y meet in A but are disjoint on Q; gets label B if they meet in B but are disjoint on Q; and label Q if they meet on Q (and therefore meet both in A and in B). We use this edge labeling to produce a vertex labeling, with the same labels as follows: if any of our sets X is the AA vertex of any AAB triangle then it receives the label A. If X is the BB vertex of any BBA triangle then it gets the label B; and otherwise it gets the label Q.

We will next show that the collection of our convex polygons which have label A form a special configuration, as do those with label B after replacing the word "above" by "below" in the definition of same. We then show that the polygons with label Q can be absorbed into one or another of these special configurations after removal of those polygons that contain either of two particular points on Q.

Recall that a special configuration of polygons consists of those obeying the two conditions:

1. Each pair of polygons meet on or above Q; each polygon meets Q.

2. Any three polygons that are mutually disjoint on Q have a point in common above Q.

We already know that every one of our polygons meets both A and B and every pair of them meet in one or the other, so that we need only verify the second condition. If we show that three polygons with label A that are mutually disjoint on Q must meet one another only on the A side of Q, from our original condition applied to them and B we can deduce that the second condition is satisfied. Thus, to get the desired conclusion here we must show that every pair consisting of two of our polygons both with label A must get label A or Q, and not B. Proof of this last claim is contained in the next subsection, in which we also observe that no vertex can receive both labels, A and B.

5 Polygons with A labels form a Special Configuration

By the remarks immediately above, this claim follows immediately from the following theorem:

Theorem 2: *Two of our convex polygons that both have A labels cannot form a pair with edge label B. Moreover, any with label A cannot also have label B.*

Proof: The definitions of labelings imply that these claims are equivalent to the impossibility of the labeled subgraphs of the complete graph illustrated in Figure 5. We show them to be impossible by invoking two principles, one logical and the other geometric.

Rule 1: In order for three sets to have a single point in common, each pair of them must meet on the same side of Q or on Q itself. Thus, by our original condition, among any four of our vertices, there must be a triangle in which the edge labels are all A or Q, or B or Q. This cannot happen if among these four vertices there is a 'matching' of B edges and one of A edges: if, for example, the edges $(1, 2)$ and $(3, 4)$ are B edges, while $(1, 3)$ and $(2, 4)$ are A edges. This is so because, among four vertices, every triangle meets every matching. Our original condition therefore prohibits simultaneous A and B matches among 4 of our polygon-vertices (see Fig. 4).

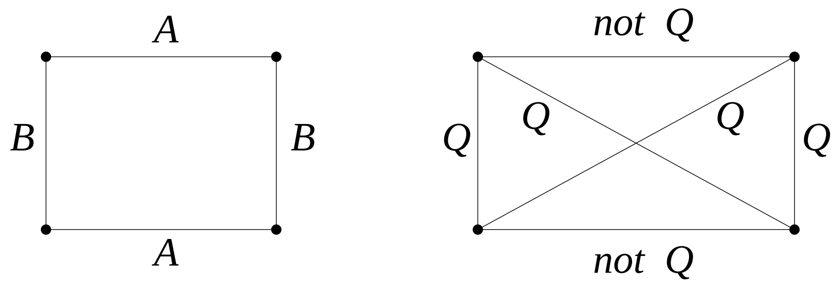

Figure 4. Forbidden configurations.

Rule 2: Suppose polygons 2 and 3 are disjoint on Q as are polygons 4 and 5. Then all four of the pairs $(2, 4)$, $(2, 5)$, $(3, 4)$, and $(3, 5)$ cannot meet on Q. For, if 2 and 3 are disjoint on Q, any polygon that meets both on Q must include the interval between them; any two such polygons must therefore themselves overlap on Q (see Fig. 4).

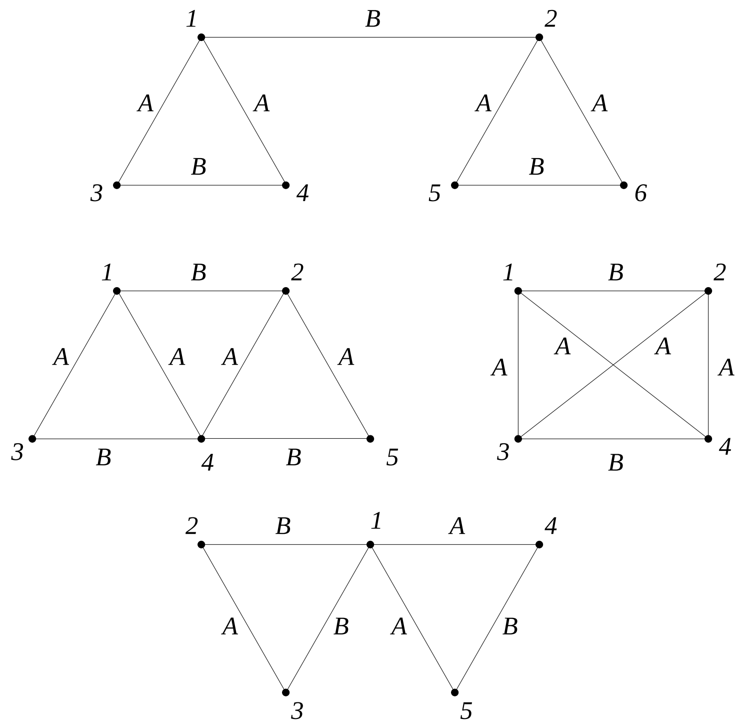

Figure 5. Configurations with B edge between A vertices, and simultaneous A, B vertices.

The first of these rules immediately eliminates the middle two subgraphs of Figure 5, when applied to the quadrilateral $(1,2,3,4)$. In the bottom case, the first rule implies that each of $(3,5)$, $(2,5)$, $(3,4)$ and $(2,4)$ meet on Q which contradicts the second rule.

In the top case, the first rule applied to $(1,2,3,4)$ implies that the edges $(2,4)$ and $(2,3)$ must each be labeled either Q or B, and the same argument applies similarly in $(1,2,5,6)$ to the edges $(1,5)$ and $(2,6)$. But if any of these are B labels, we have a bottom type configuration which we have already ruled out. We may conclude that all of (the four edges, $(1,5)$, $(1,6)$, $(2,3)$ and $(2,4)$ have Q labels.

The first rule applied to $(1,2,3,5)$ implies that the edge $(3,5)$ has label A or Q; and symmetrically, the same conclusion follows for edges $(3,6)$, $(4,5)$ and $(4,6)$. The second rule then, applied to $(1,3,5,6)$ implies that one of $(3,5)$ and $(3,6)$ must have label A. By applying this argument symmetrically we find that each of 3, 4, 5 and 6, must have an A labeled edge to another; this implies that there is an A labeled matching which gives us a contradiction with our first rule: there is already a B labeled matching among these vertices.

So the A and B labeled polygons each form special configurations. We now consider Q labeled polygons.

6 Q labeled Polygons

Let L be the leftmost Q labeled polygon on Q; that is, the one whose right endpoint on Q is furthest to the left; let R be similarly, the rightmost Q labeled polygon on Q. If L and R meet on Q, then all Q labeled polygons meet, and there is a single point, p' on Q that meets them all, by Helly's theorem in one dimension.

We therefore assume otherwise, that L and R are disjoint on Q and the pair (L, R) has label A or B. Suppose, without loss of generality that this label is A.

Here are two facts that give us two points that allow us to merge Q and A among the polygons missing them.

Fact 1: Three Q labeled polygons that do not meet one another on Q must have all of their edges A's or all of their edges B's.

Proof: Any mixed triangle has a vertex with both labels the same which vertex would be an A or a B rather than a Q labeled polygon.

Fact 2: If two Q labeled polygons are disjoint from each other and from an A labeled polygon on Q, and if they form an A labeled pair, then all edges of their triangle must be labeled A's.

Proof: The triangle cannot be edge labeled ABB or else the A labeled vertex would also have a B label, which contradicts Theorem 2. If the edge labels were AAB one of the two Q labeled vertices would have an A label instead.

We now show that if we omit all Q labeled polygons that contain either the rightmost point, rL, of L on Q, or the leftmost point, lR, of R, then the rest of the members of Q form only Q or A labeled pairs with each other or with A labeled polygons.

Let Y be any Q labeled polygon not containing either rL or lR. Applying the first fact to the triangle (L, R, Y) we find that (L, Y) and (R, Y) must have A labels. If Y and Y' have Q labels and miss lR and rL, they either meet on Q and (Y, Y') has label Q, or, applying the first fact to either triangle (L, Y, Y') or (Y, Y', R) we find that (Y, Y') must have label A. Finally, if X has A label and lies to the left of Y on Q and is disjoint from Y on Q, then applying the second fact to the triangle (X, Y, R), noting that (Y, R) has label A, we deduce that (Y, X) has label A, which completes proof of the claim of the previous paragraph. The same argument applies if A lies to the right of Y on Q with L and R interchanged. We can summarize this discussion by the statement that all (X, Y) or (Y, Y') edges are necessarily labeled Q or A.

These statements assure us that we can absorb all the polygons of Q into A without introducing any B labeled edges among them and therefore maintaining the resulting collection as a special configuration, after removing those polygons containing the explicit two points, the right end, rL, of L and the left end, lR, of R, which points we choose as p' and p''.

Consequently, the sets that do not contain any of p, p' and p'' can be divided into two special configurations, both of which can be pierced by five points. So, with 13 points we can pierce all sets.

7 Comments on The Argument

The conclusion we draw from the arguments presented above is that 13 points suffice to cover all the polygons. We used three points to set up the special configurations of A's and B's and five to cover them independently. If you study the problem you will soon be convinced that no more than three points are required to handle a special configuration, and probably to handle any set of polygons obeying the given conditions. Though this problem has been open for more than forty years, it is probable that this particular case was not studied very hard, because of its special nature. It may well be that there is a simple argument for a tighter bound that is lurking, waiting to be found. We can attest to the existence of a goodly number of false proofs of such bounds. An interesting open question is: do these arguments help at all in any other cases of this problem? Can one say anything intelligent about the case in which three out of every five have a point in common, among closed convex sets in the plane? Or four out of every five closed convex polyhedra do so in three dimensions?

In the spirit of Paul Erdős, the authors will give $x \cdot \$10$ for any improvement of the upper bound by x, below 13, and \$30 for each incremental improvement of the lower bound above 3 for the problem considered here. Again in the Erdős spirit, the argument presented above possesses a certain esthetic charm. It will be a pity if and when it is replaced by a still more direct and more conventional argument which gives a better bound. It is not clear how much effort has gone into attempting to improve the lower bound for this problem. The lower bound three can not be improved for special configurations, this is shown by the following example with six triangles. Select a regular triangle T with vertices A_i and let B_i be the point which divides the segment A_iA_{i+1} in ratio $1:2$ and closer to A_i $(i \in \{1,2,3\})$. Now the special configuration is defined by the three regular triangles whose sides are the sides of T and their third vertex is outside T together with the three triangles A_i, B_i, B_{i+1}. It is easy to check that the six triangles form a special configuration with respect to the line A_1A_2, three of any four meet and no two points cover all of them. However, a configuration requiring more than three points for a cover in the general case has not been found.

References

[ABFK92] N. Alon, I. Bárány, Z. Füredi and D. J. Kleitman, Point selections and weak ε-nets for convex hulls, *Combinatorics, Probability and Computing* **1** (1992), 189–200.

[HD57] H. Hadwiger and H. Debrunner, Über eine variante zum Hellyschen satz, *Arch. Math.* **8** (1957) 309–313.

[AK92] N. Alon and D. J. Kleitman, Piercing convex sets and the Hadwiger–Debrunner (p,q) problem, *Advances in Mathematics* **96** (1992), 103–112.

[AK97] N. Alon and D. J. Kleitman, A purely combinatorial proof of the Hadwiger–Debrunner (p,q) conjecture, The Wilf Festscrift, Philadelphia 1996, *Electron. J. Combin.* **4** (1997), no. 2, Research paper 1.

[GL70] A. Gyárfás and J. Lehel, A Helly-type problem in trees, in: *Combinatorial Theory and its Applications, Eds. P. Erdős, A. Rényi and V. T. Sós,* North Holland, Amsterdam 1970, 571–584.

[GL85] A. Gyárfás and J. Lehel, Covering and Coloring problems for relatives of intervals, *Discrete Mathematics* **55** (1985), 167–180.

[KL79] M. Katchalski and A. Liu, A problem of geometry in R^d, *Proc. Amer. MAth. Soc.* **75** (1979), 284–288.

[W75] G. Wegner, d-collapsing and nerves of families of convex sets, *Arch. Math.* **26** (1975), 317–321.