Discrete Geometry: Independent Study Report

Athreya Chandramouli (2018121002)

May 2020

1 Work done as part of Independent Study

1. The aim of the study was to cover the sections marked as "Foundations: basic things suitable for an introductory course" from Lectures on Discrete Geometry by Jiri Matousek [1].
2. This includes the following sections-
   - 1.1-1.3: Linear and Affine Spaces, Convex Sets, Combinations, Separation.
   - 2.1: Minkowski's Theorem.
   - 5.1-5.4, 5.7: Geometric Duality, Convex Polytopes, Voronoi Diagrams.
3. Apart from these I studied the following from lecture notes by Prof. Rajeev Raman for the ACM Winter School on Geometric Algorithms.
   - Crossing Numbers in graphs
   - Point-Line incidences, Szemeredi-Trotter
   - Clarkson-Shor technique
   - Polygon Triangulation

The following sections attempt to cover a lot of the major concepts succinctly. Many non-trivial statements are stated without proofs in the interest of keeping the report brief. I have however worked through most of the proofs and have attempted some of the exercises in the textbook as a part of the study. Some interesting results from the exercises are also mentioned inline.

2 Convexity

2.1 Affine Subspaces

Affine subspaces serve as the natural extension to the notion of a linear subspace. While linear subspaces are required to pass through the origin (for ex: lines, planes, hyperplanes through origin), affine subspaces are not bound by this restriction. Thus a way to define an affine subspace of $\mathbb{R}^d$ is a subspace of the form $x + L$ where $x \in \mathbb{R}^d$ and L is a linear subspace of $\mathbb{R}^d$.

Using this above definition, we can then extend linear notions such as as Linear Span, Linear Combinations etc.

An Affine Span, more commonly known as an Affine Hull, of a set of points X is defined to be the intersection of all affine subspaces of $\mathbb{R}^d$ containing X.

The affine combination of points $a_1, a_2, \cdots, a_n \in \mathbb{R}^d$ is an expression of the form $\alpha_1 a_1 + \alpha_2 a_2 + \cdots + \alpha_n a_n$ where $\forall i \alpha_i \in \mathbb{R}^d$ and $\sum_i \alpha_i = 1$.

We can extend the notion of linear independence as well. A collection of points is said to be affinely dependent if one of them can be written as an affine combination of the others. This implies the existence of $\alpha_1, \cdots, \alpha_n$ such that at least one of them is non-zero such that they satisfy both $\sum_i \alpha_i a_i = 0$ and $\sum_i \alpha_i = 0$. Playing with the terms above we can get that a set of n affinely dependent points $a_1, \cdots a_n$ in $\mathbb{R}^d$ corresponds to a $n-1$ linearly dependent points $a_1 - a_n, \cdots a_{n-1} - a_n$. Since there can only be at most d linearly independent points, we can conclude there can be at most $d+1$ affinely independent points. This observation also leads us to an efficient way to check affine dependence: construct a matrix where each column corresponds to $a_i - a_n$. If its determinant is non-zero, then the points are affinely dependent. In order to refer to subspaces succinctly, we define a k-dimensional affine subspace as a k-flat.

The notion of a linear transform extends to an affine mapping: a function of the form $f : y \mapsto By + c$ where B is a linear transform represented by a $d \times k$ matrix and $c \in \mathbb{R}^d$.

A common idea that appears throughout many proofs is that of general position. At a high level this simply means that there are no unlikely coincidences. For example, if 3 points are chosen with no special intention, they're unlikely to lie on a line. The precise meaning is not fully standard and depends on the context in which it is used.

2.2 Convex Notions

A set C is said to be convex if for every two points $x, y \in C$, the whole segment xy is contained in C. i.e. for every $t \in [0, 1]$, the point $tx + (1-t)y \in C$.

The convex hull of a collection of points is defined as the intersection of all convex sets containing the points. This is denoted as $conv(X)$.

A convex combination is defined as $\sum_i t_i x_i$ for $t_i \in [0,1]$ and $\sum_i t_i = 1$.

An interesting result about convex hulls is that the entire $conv(X)$ can be written as the union of convex combinations of at most $d+1$ points in X. For example, any convex set in a plane can be written as the union of all triangles with vertices at points of X. This is Caratheodory's Theorem.

The separation theorem suggests that two distinct (having no intersection) convex sets can always be separated by a hyperplane.

Farkas Lemma suggests a connection to Linear programming. It states that for every $d \times n$ real matrix A, exactly one of the following occurs: a) The system has a non-trivial, non-negative solution. or b) There is a vector y such that $y^T A$ is a vector with all entries strictly negative. This is a special case of the duality of linear programming.

2.3 Radon's lemma and Helly's theorem

Radon's lemma, Caratheodory's theorem, and Helly's theorem form the three basic properties of convexity in $\mathbb{R}^d$.

Radon's lemma suggests that for a set A of $d+2$ points in $\mathbb{R}^d$, there are two disjoint subsets $A_1, A_2 \subset A$ such that $conv(A_1) \cap conv(A_2) \neq \phi$.

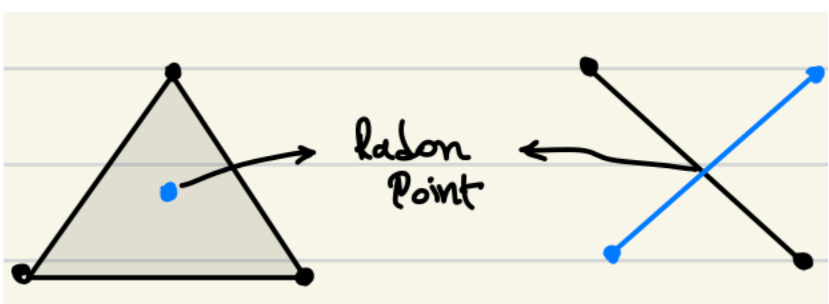

.

For example, in the plane, $d = 2$, so for $d+2 = 4$ points in 2D here are two possible partitions shown in blue and black. Points in the intersection of their convex hulls are called Radon Points.

Helly's theorem is an interesting combinatorial result. It states that for convex sets $C_1, \cdots C_n$ in $\mathbb{R}^d$, where $n \geq d+1$, if the intersection of every d+1 of these sets is nonempty, then the intersection of all of them is nonempty. The theorem is useful since in proofs it allows us to deal with only the intersection of $d+1$ convex sets instead of an arbitrary number. This proved very useful in attempting the exercises in the textbook where it was asked to prove that a compact convex set of width 1 contains a segment of length 1 in every direction, among other problems.

3 Lattices and Minkowski Theorem

For any basis of $\mathbb{R}^d$ the subgroup of all linear combinations with integer coefficients of the basis vectors forms the integer lattice $\mathbb{Z}^d$.

Minkowski's theorem states that for a convex, symmetric(around the origin), bounded such that $vol(C) > 2^d$. Then C contains at least one lattice point different from the origin. The proof follows from defining a set $C' = \{\frac{1}{2}x | x \in C\}$ and showing that C' and one of its integer translates intersect.

This theorem allows us to reason about problem like: given a circle K of some diameter with center at origin, and that smaller circles of small diameter grow at every lattice point, is there a line from the origin to a point outside the circle that doesn't intersect any of the smaller circles. Here if we reframe the problem as whether the strip of width equal to the diameter of the smaller circles contains one of the lattice points, all we need to check is whether the volume condition is met. Another interesting application is to find the fraction approximation of irrationals, where we can reason about the strip of width equal to the allowed error in approximating around the line representing the irrational and reason that since this is a convex space satisfying the conditions of Minkowski's theorem, we can find a lattice point which in turn gives us a fractional representation.

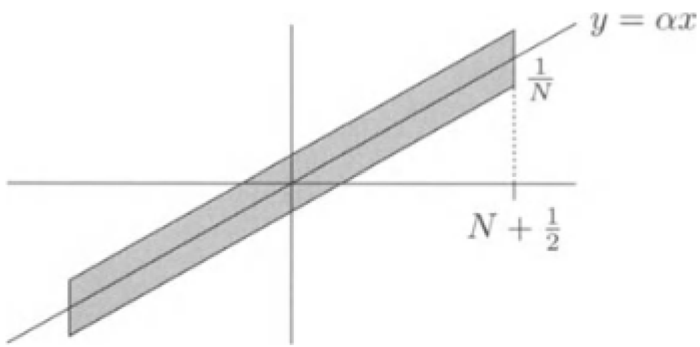

.

Here we see that if we want the error to be within $\frac{1}{N}$ we can consider the strip $C = \{(x, y) | -N - 0.5 \leq x \leq N + 0.5, |\alpha x - y| < \frac{1}{N}\}$.

4 Geometric Duality

The geometric duality transform is a mapping s.t. for a point $a \in \mathbb{R}^d - \{0\}$ it assigns the hyperplane: $D_0(a) = \{x | a^T x = 1\}$. and vice versa. Geometrically, a point at a distance d from the origin maps to a hyperplane perpendicular to the line joining the point and the origin at a distance of $\frac{1}{d}$ from the origin.

Another interpretation is to move to $\mathbb{R}^{d+1}$ and define the primal hyperplane as the one with $x_{d+1} = 1$ and the dual as that with $x_{d+1} = -1$. Now to find the dual hyperplane, in $\mathbb{R}^{d-1}$ we construct the hyperplane passing through origin that is perpendicular to the line connecting the origin and the point on the primal plane. The intersection of this hyperplane with the dual hyperplane is the required dual for the point.

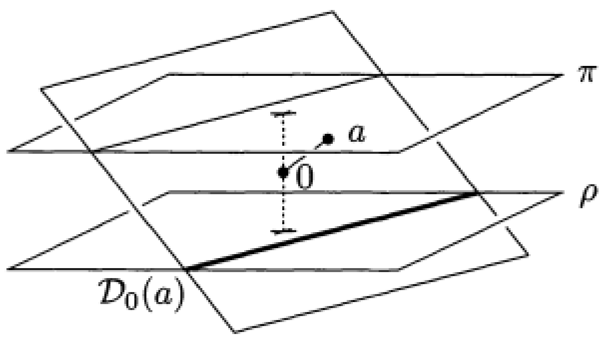

.

This allows us to extend the transform to k-flats in $\mathbb{R}^d$ to give (d-k-1)-flats. This is done by finding a hyperplane passing through origin that contains the k-flat. The intersection of the orthogonal complement of this hyperplane with the dual hyperplane is the required dual.

This transform is useful since it preserves incidences: i) $p \in h \Leftrightarrow D_0(h) \in p$ and ii) $p \in h^- \Leftrightarrow D_0(h) \in D_0(p)^-$ where h^- represents the half-space bounded by h.

We can thus also define a Dual Set for a set of points X as $X^* = \{y | x^T y \leq 1 \ \forall x \in X\}$. Geometrically this is the intersection of all half-spaces of the form $D_0(x)^-$ with $x \in X$. Its interesting to note that fora closed convex set containing the origin $(X^*)^* = X$.

5 Polytopes

There are two equivalent definitions of polytopes. An H-polyhedron is defined as an intersection of finitely many closed half-spaces in $\mathbb{R}^d$. An H-polytope is a bounded H-polyhedron.

Alternatively, a V-polytope is the convex hull of a finite point set in $\mathbb{R}^d$.

The equivalence can be proved by inducting on the number of dimensions. However the two formulations can result in drastically different computations. For example, consider the problem of optimization. In V-polytopes, it is sufficient to simply evaluate and compare over the points in the set, however for H-polytopes, it requires solving a linear programming problem which we know is non-trivial.

Certain particular polytopes are-

1. **Permutahedron**: A d-dimensional polytope given by the convex hull of the $(d+1)!$ vectors in $\mathbb{R}^{d+1}$ that arise by permuting $[d+1]$. These always like on a d-dimensional subspace since $\mathbf{1}^T X = \frac{n(n+1)}{2}$

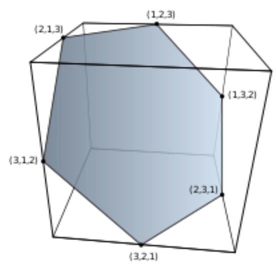

2. **Cubes**: Defined by the point set $[-1,1]^d$. This can of course be represented as a H-Polytope $-1 \leq x_i \leq 1$. Similarly its V-polytope representation is the convex hull of $\{-1,1\}^d$.

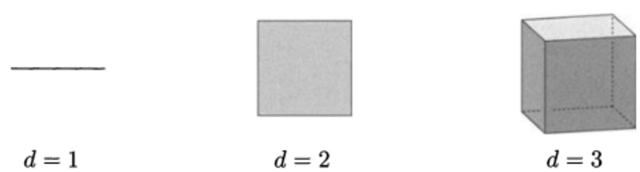

3. **Crosspolytopes**: Defined as a V-polytope as the convex hull of $\{-e_1, e_1, -e_2, e_2, \cdots, -e_d, e_d\}$ where the e_i are the vectors of the standard orthonormal basis.

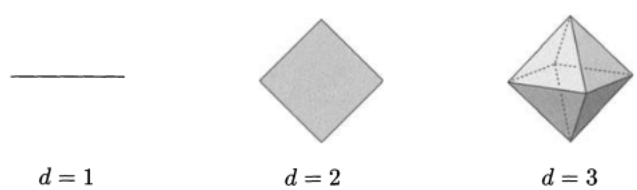

4. **Simplex**: Defined as the convex hull of an affinely independent point set in $\mathbb{R}^d$. Recall that there can be at most $d+1$ such points. A regular simplex is one with all $d+1$ points with all pairs of points being equidistant. A neat

coordinate representation is $\{e_1, e_2, \cdots, e_{d+1}\}$ where e_i are the vectors of the standard orthonormal basis in $\mathbb{R}^{d+1}$.

$d = 0$ $d = 1$ $d = 2$ $d = 3$

5.1 Faces

The face of a convex polytope P is defined as either P itself or a subset of the form $P \cap h$ where h is a hyperplane such that P is fully contained in one of the closed half spaces determined by h. For example in a 3D Cube this would be the cube itself, the square faces, the edges, the vertices and the empty face.

Notice that each face is in itself a convex polytope since it is defined by the intersection of the hyperplane h with the hyperplanes that define P. Faces can have dimensions from -1(empty face) to d(P itself). A J-face is a face of dimension J.

We can also get a graph representation of a polytope: The vertices of the graph are the vertices(0-faces) of the polytope. There is an edge between two vertices in the graph if there is an edge(1-face) containing both the 0-faces corresponding to each of the vertices. It can also be shown that 3D polytopes have Planar Graph representations. This can be constructed by projecting onto the circumscribing sphere and project the sphere onto cartesisan coordinates by a stereographic projection. This graph is also 3-connected since the simplex is 3-connected.

Steinitz Theorem states that a finite graph is isomorphic to the graph of a 3D convex polytope iff its planar and 3-connected. However for a general polytope, it is likely NP-Hard to check whether it is isomorphic to a graph.

There is also Balinski's theorem that states that the graph of a d-dimensional polytope is d-connected. But this only shows that its a necessary condition and not that it is sufficient.

Let us count the faces of the example polytopes.

1. **Regular Simplex**: Every subset of points define a face. Thus there are $\binom{d+1}{J}$ J-faces and 2^{d+1} faces overall. We can verify this for a tetrahedron: 3 triangles, 6 edges, 4 vertices, 1 empty face and the tetrahedron itself.

2. **Crosspolytopes**: A subset determines a face iff i) $\forall i\ e_i \in F \rightarrow -e_i \notin F$ and ii) $\forall i!; -e_i \in F \rightarrow e_i \notin F$. Thus there are 3 options for each i, either choose e_i, $-e_i$, or neither.Thus there are $3^d + 1$ faces in total. We can

verify this for an octahedron: 8 triangles, 12 edges, 6 vertices, 1 empty face and the octahedron itself.

3. **Cubes**: There are 3^d vectors representing the centre of each face given by $\{-1, 0, 1\}^d$. The points that make up the face is given by $\{u | u_i = v_i \; \forall \; v_i \neq 0\}$. Thus here to there is a total of $3^d + 1$ faces. For a 3D cube we can verify that there are 6 squares, 12 edges, 6 vertices, 1 empty face, and the cube itself.

5.2 Face Lattice

Let $F(P)$ be the set of faces for a bounded convex polytope P. Two polytopes P and Q are said to be combinatorially equivalent if $F(P)$ is isomorphic to $F(Q)$ as Posets. Here a face is comparable to and smaller than an other face if it is contained in the other face.

In fact $F(P)$ is actually a lattice. i.e. it is a poset satisfying-

1. Meets Condition: For any two faces F and G of P, there exists a face M that is contained in both F and G and contains all other faces contained by both F and G.

2. Joins Condition: For any two faces F and G of P, there exists a face J that contains both F and G and is contained in all other faces that contain both F and G.

This lattice has the following properties-

1. Graded: Every maximal chain is of the same length.

2. Atomic: Every face is the join of its vertices.

3. Co-Atomic: Every face is the meet of the facets that contain it.

This makes it a good representation to store the polytope computationally, since we just have to store the face the facets containing it and the faces it contains. This is however not used in practice since for most applications a much sparser representation suffices.

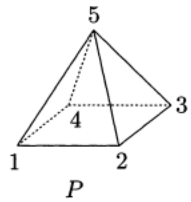

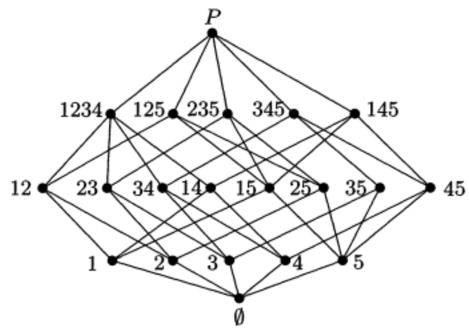

We defined the dual set of a set of points. The dual set of a polytope is also a polytope known as the dual polytope. Each of the j-faces of P are in bijective correspondance with the $(d-j-1)$-faces of P^*. It is also worth noting that the face lattice of P^* arises by turning the face lattice of P upside down.

We define a simplical polytope as one where each of its facets is a simplex. A simple polytope is one where each vertex is contained in exactly d facets. Notice that crosspolytopes are both simplical and simple. For a counter example, a square pyramid is neither simplical nor simple. The dual of a simplical polytope is simple and vice-versa.

5.3 Cyclic Polytopes

Let f_i be the no. of i-faces in the polytope P. Now given the n vertices of the polytope, what is the max total no. of faces that it could have? For $d = 2$ it is apparent that there are n vertices and n edges. For $d = 3$ we know that there is a planar graph corresponding to the polytope, hence $f_1 \leq 3n - 6$ and $f_2 \leq 2n - 4$ with equality when the polytope is simplical. In these cases the #edges is linear in n. However as dimensions grow, even for a simplex, #faces is exponential in d.

A moment curve is the curve $\gamma = \{t, t^2, \cdots, t^d | t \in \mathbb{R}\}$. Now any hyperplane will intersect this curve at at most d points. At each of these points the curve changes which side of the hyperplane it is on.

A cyclic polytope is defined as the convex hull of finitely many points on the moment curve. Here Gale's evenness criterion states that for a vertex set V considered with a linear ordering along the moment curve, then any d-tuple of the vertices in the polytope determines a facet iff for any two vertices u,v not in the d-tuple, the no. of vertices in the d-tuple between u and v is even. Using this it can be proved that the #facets of a d-dimensional cyclic polytope with n vertices is $\binom{n-\lfloor d/2\rfloor}{\lfloor d/2\rfloor} + \binom{n-\lfloor d/2\rfloor-1}{\lfloor d/2\rfloor-1}$ for even d and $2\times \binom{n-\lfloor d/2\rfloor-1}{\lfloor d/2\rfloor}$ for odd d.

6 Voronoi Diagram

For a finite set of points P, let reg be a region defined as all points for which a point $p \in P$ is the closest among the poins in P. More formally, $reg(p) = \{x \in \mathbb{R}^d | dist(x, q) \leq dist(x, p) \; \forall \; q \in P\}$.

The Voronoi diagram of P is the set of all regions $reg(p)$ for $p \in P$.

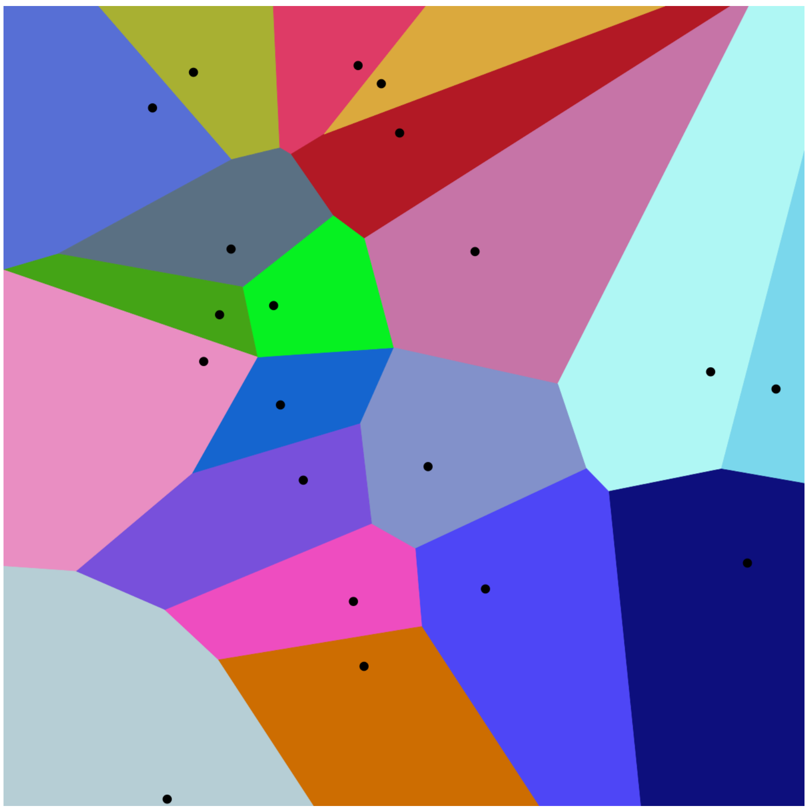

Notice that each region is a convex polyhedron with at most $|P| - 1$ facets.

Some potential applications are-

1. Nearest neighbour search: given a point, finding the Voronoi region the point belongs to gives us the nearest neighbour of that point.
2. Robot Motion Planning: given a set of point obstacles and a disk shaped robot that has to reach from start to finish without touching these points. The robot can always follow the edges of the Voronoi diagram. Thus the problem becomes a graph search which is well studied.
3. Delaunay Triangulation: Many applications require that a convex region

is divided into triangles. This is known as the triangulation problem. A desirable properties are that not all triangles are skinny among others. The Delaunay triangulation connects an edge between two points in P if their Voronoi regions share an edge. This turns out to be a good triangulation with respect to several criterions.

4. Interpolation: Given some finite smooth function for $R^2 \rightarrow R$, who values are know to us at a finite set of points P, we can interpolate the value for some $x \in conv(P)$ as follows. Construct the Voronoi diagrams for both P and $P \cup \{x\}$. Now let w_p be the area of the part of $reg(p)$ that belongs to $reg(x)$ after inserting x. Then we can interpolate $f(x)$ as $\sum_{p \in P} \frac{w_p}{\sum_{q \in P} w_q} f(p)$

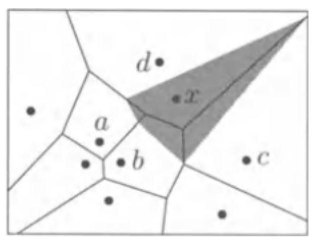

7 Crossing Number of Graphs

A drawing of a graph is a mapping of the vertices to $\mathbb{R}^2$ and the edges to continuous curves between the vertices.

The crossing number of a graph $cr(G)$ is the smallest no. of crossings of edges in any drawing of G.

A planar graph is a graph that can be drawn in the plane with no crossings.

Now, Kuratowski's theorem states that a graph is planar iff we cannot obtain K_5 (complete graph with 5 vertices) or $K_{3,3}$ (complete bipartite graph with 3 vertices in each partition) via a sequence of edge contractions and vertex deletions.

Here an edge contraction replaces two end points with a single vertex whose neighborhood is the union of neighborhoods of the two end points.

We are interested in bounds on the crossing number of various graphs. For example we can show that $Cr(K_n) = \omega(n^4)$ using a double counting argument by counting the contributions of an induced subgraph K_5.

We can also argue that for any graph G with m edges and n vertices, $Cr(G) \geq m - 3n$. We can argue this by removing one edge for each crossing (thus removing $Cr(G)$ edges). This leaves us with a planar graph which we know has at most $3n-6$ edges. Hence $m - Cr(G) \leq 3n-6$ and playing around with these terms gives us the required bound.

We can also show that $Cr(G) = \omega(m^3/n^2)$ which is proved by a similar double counting argument that uses size t induced sub graphs instead of K_5. However a much more direct proof follows from Szekely who gave a neat probablistic method proof.

Let $p \in (0,1)$ be determined later. Consider an induced subgraph of G, say G' formed by selecting vertices independently with probability p.

Now $Cr(G') \geq |E(G')| - 3|V(G')|$. Now by linearity of expectation the expected size of $V(G') = np$, $E(G') = mp^2$ and so expected no. of crossings is $Cr(G') = Cr(g)p^4$ (all 4 vertices must be chosen).

Thus in expectation $Cr(G)p^4 \geq mp^2 - 3np$. If we choose $p = 4n/m$ we get $Cr(G) \geq \frac{1}{64}\frac{m^3}{n^2}$.

8 Point-Line Incidences

Let P be a set of points and L be a set of lines. Now we define $I(m,n)$ be the number of pairs (l,p) such that $p \in l$ for $l \in L$ and $p \in P$.

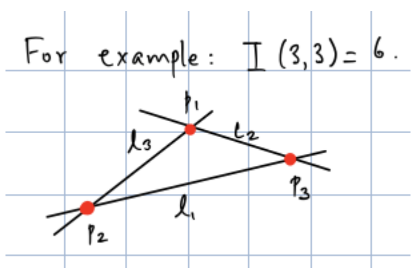

A trivial bound is obviously $I(m,n) \leq mn$.

Szemeredi and Trotter proved that the #incidences among n lines and m points is at most $O(m^{2/3}n^{2/3} + m + n)$. The proof uses the bound on crossing numbers.

The Szemeredi Trotter bound can be used for several other problems. For example: Erdos posed a question to determine the maximum no. of pairs of points in the plane that can be at distance 1. This bound turns out to be $O(n^{4/3})$. This is acheived by modelling the problem as given n unit circles how many centres lie on another circle's boundary. Thus a pair of points at unit distance contributes two incidences with unit circles. To count the #incidences we can use a similar argument to Szemeredi Trotter to show that it is indeed $O(n^{4/3})$.

9 Clarkson-Shor Technique

Consider a motivating problem. Given a set of lines in the plane, they determine a set of points: the points of intersection of each pair of lines.

The depth of a point is the number of lines lying below it.

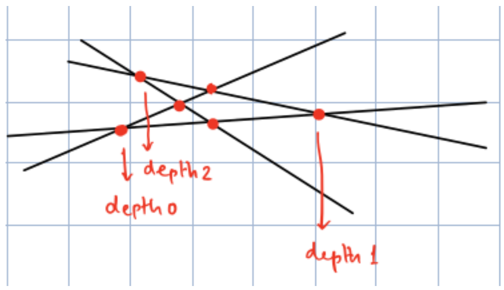

We want to find the maximum no. of points of depth at most k in an arrangements of n lines. We cam prove that this no. is in face $O(nk)$.

Let L be any set of n lines. We pick a sample R from L where each line is chosen with probability p independently.

Now we can observer that the #points at depth 0 in an arrangement of n lines is at most $n-1$. The expected size of R is np. Now the probabilty that a point at depth $\leq k$ is at depth 0 in R can be calculated as P(both lines defining x are chosen in R and no line below x is chosen in R). This is $p^2(1-p)^j$, where j is the depth of x in L. This is $\geq p^2(1-p)^k \geq p^2e^{-2pk}$. For $p = 1/k$, the above probability is $\geq \frac{1}{e^2k^2}$. Thus the expected #points at level 0 in R is at least $\frac{|V_{\leq l}|}{e^2k^2}$.

Thus in expectation $\frac{|V_{\leq l}|}{e^2k^2} \leq |R| \leq \frac{n}{k}$ which gives us our required bound.

Consider a more general version of this problem. Let L be a set of n objects and ρ be a set of configurations where each configuration is defined by $\leq d$ objects. Given a conflict relation $R \subset L \times \rho$, we are interested in the number of configurations with $\leq k$ conflicts.

Clarkson and Shor proved that $N_{\leq k}(L) \leq \frac{\mathbb{E}[N_0(R_p)]}{p^d(1-p)^k}$ where $N_0()$ is the no. of objects with no conflicts.

For example: let P be a set of n points in $\mathbb{R}^2$. Now define a graph $G_k(P)$ where there is an edge p_ip_j iff there is a disk containing p_i and p_j and $\leq k$ other points.

It is easy to see that $G_0(P)$ is a planar graph. We can also see that the #edges in a random subset of P is at most $3np$. Similarly, the expected size of the vertices is found to be $mp^2(1-p)^k$. For $p = \frac{1}{2k}$ we eventually get $m \leq \frac{2nk}{e^4}$.

Another example: let P be a set of n points in the positive quadrant of $\mathbb{R}^2$. Let R be a set of rectangles anchored by the x-axis and having 3 points on each side. The number of such rectangles that have $\leq k$ points in their interior is $O(nk^2)$. This once again follows using a similar argument and applying Clarkson-Shor.

10 Polygon Triangulation

A triangulation is a decomposition of a polygon P into triangles by a maximal set of non-intersecting diagonals.

It is in fact non-trivial to even prove that a diagonal exists. A simple visual proof follows from these two figures. Consider the vertex v and its two adjacent vertices u and w.

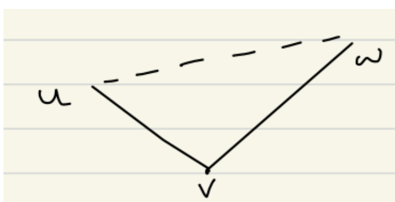

If the line segment connecting uw lies entirely inside the polygon we trivially find a diagonal. Otherwise, consider a sweep line parallel to uw moving upwards from v. The first vertex that the line intersects, say x results in the existence of a diagonal vx.

Knowing this, we can show that every polygon has a triangulation by inducing on the no. of vertices (a diagonal splits the polygon into two smaller polygons).

We define P to be a monotone polygon with respect to the y-axis if every intersection of a line perpendicular to the y-axis with P is connected.

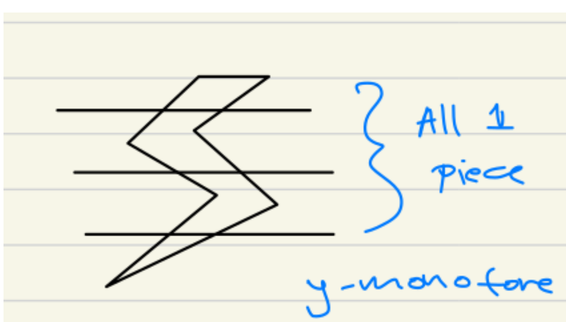

A computational way to find a triangulation would be to first divide the given polygon into monotone polygons and then triangulate those monotone polygons. In fact this can be done quite efficiently. Chazelle discovered an algorithm that actually runs in $O(n)$ time.

It can be proved that every polygon triangulation in 2D has $n-2$ triangles.

Interestingly, even if the polygon has holes, it still admits a triangulation. Unlike earlier however we cannot induct since a diagonal need not divide the polygon into two smaller parts. However now we can notice that drawing a diagonal to a hole eliminates the hole. This we can induct on #holes and #vertices.

Further, it can be proved that for a polygon with h holes and n vertices the triangulation has $n+2h-2$ triangles.

A 3D generalization of polygon triangulation is to decide whether a polytope can be deomposed into tetrahedrons formed using the vertices(0-faces) of P. However, even deciding whether such a decomposition exists is an NP-Complete Problem.

Another interesting fact is that the #triangulations is in fact the Catalan Number.

10.1 Art Gallery Problem

Given a polygon without holes, place guards such that the entire polygon is visible. We can show a nice lowerbound using this example.

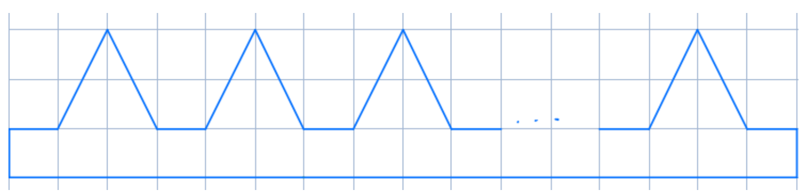

Here we can see a polygon with $\frac{n}{3}$ triangular peaks. Here $\lfloor \frac{n}{3} \rfloor$ guards are necessary.

An easy upperbound is to notice that we can simply triangulate the polygone and place one guard inside each triangle. This would require $n-2$ guards.

Note however, that if we can 3-color vertices of the polygon, then there must be at least one color class of size $\leq \frac{n}{3}$. If we place a guard at each such vertex, we can cover every triangle.

Now, an ear is a triangle such that two of its sides lie on the polygon. We can show that for $n>3$ any triangulation has at least 2 ears. This follows from the fact that there are $n-2$ triangles and n vertices.

This gives us a way to induct on the triangulation. We choose an ear and remove it. By the inductive hypothesis the smaller polygon(without the ear) can be 3-colored. However, the ear that was removed can be reinstated by coloring the vertex with the extra color not being utilized by its neighbours.

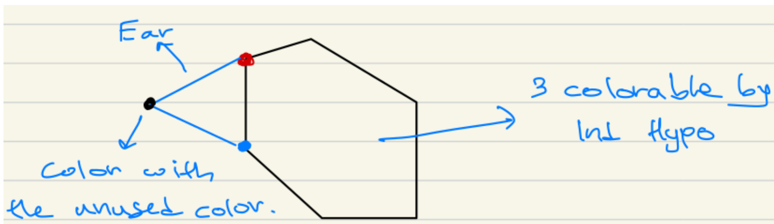

Thus a tight bound for the #guards required is $O(\lfloor \frac{n}{3} \rfloor)$.

References

[1] Jiří Matoušek, editor. *Lectures on Discrete Geometry.* Springer New York, 2002.