Reflections on convexity, connectedness and visibility: Rich unions and elegant intersections.

Sudebkumar Prasant Pal
Department of Computer Science and Engineering
I I T Kharagpur 721302
India
Email: spp@cse.iitkgp.ernet.in
http://www.angelfire.com/or/sudebkumar
http://www.facweb.iitkgp.ernet.in/~spp/

The study of shapes and topologies of objects in two- and three-dimensional Euclidean space have influenced and absorbed mathematicians for centuries. The most important and pervasive notion in shapes is that of "convexity". An object A is said to be convex provided the join pq of any two points p and q of A lies entirely in A. In Figure 1, A, B and C are convex polygons but the union of B and C is non-convex. For p in B and q in C, the join pq goes outside the union of B and C.

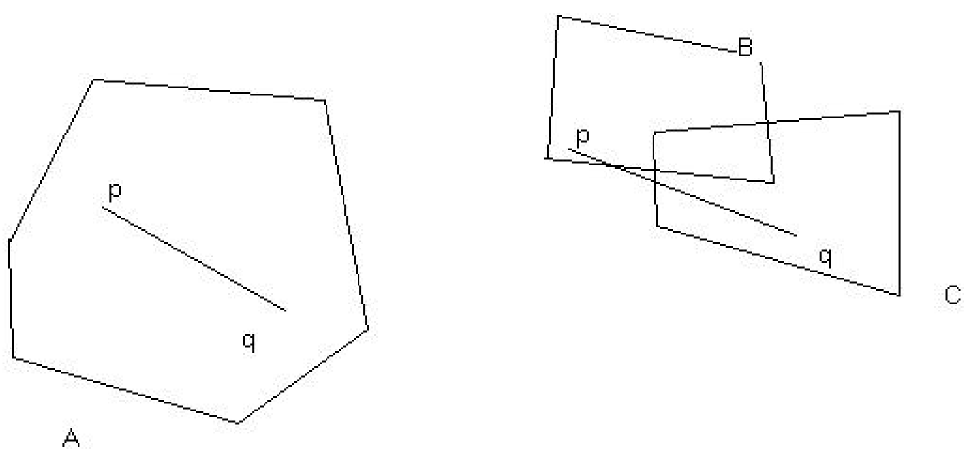

Figure 1: Convex and non-convex objects

A more algebraic characterization of convexity is as follows. An object C is said to be "convex" if for any two points (position vectors) p and q in C, ap+bq also lies in C, where a and b are positive reals and a+b=1.

It is interesting to note that the union of convex objects can be very rich indeed. It is not difficult to show that the intersection of convex objects is convex. On the other hand, the union of convex objects may be non-convex and may also have islands or holes as we see in Figure 2.

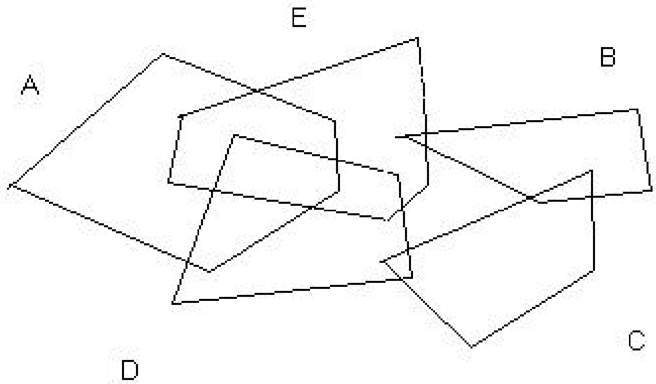

Figure 2: The common intersection of A, D and E is convex. The union of A, B, C, D and E is non-convex and has a hole or island.

Complex non-convex objects are often viewed as unions of convex objects. A garland is made of spherical beads. The problem therefore of decomposing a non-convex object into convex objects is an important computational problem, required in many applications such as pattern recognition. It is interesting to know the minimum number k of convex polygons whose union is a given non-convex polygon P. The uniting convex polygons in this case might have overlaps. We may also be interested in partitioning the same non-convex polygon P into a minimum number k' of mutually non-overlapping convex polygons. Is k=k'? Can you find a non-convex polygon P such that k is not the same as k'? The union of the five convex polygons in Figure 2 is not the union of any four convex polygons.

Unions that do not involve holes are simpler to deal with. Simple polygons are non-convex in general and may be viewed as unions of convex polygons. A simple polygon partitions the Euclidean plane into three parts: the interior (which can be shrunk to a point), the exterior, and, the polygonal boundary itself. The boundary of a simple polygon is a closed Jordan curve. An n-vertex or n-sided simple polygon is called an n-gon. It would be interesting to characterize the number of distinct shapes you could get for n-gons. For n=3, there is only one shape, that of a triangle. For n=4, you can get only two shapes as shown in Figure 3. One is a convex quadrilateral and the other is non-convex. How many shapes can you get for n=5? Can you find the function f(n) where f(n) is the number of possible shapes for n-gons? Have we defined what notion of shape we are really interested in? Ponder and define your own problems; there may not be a single elegant definition!

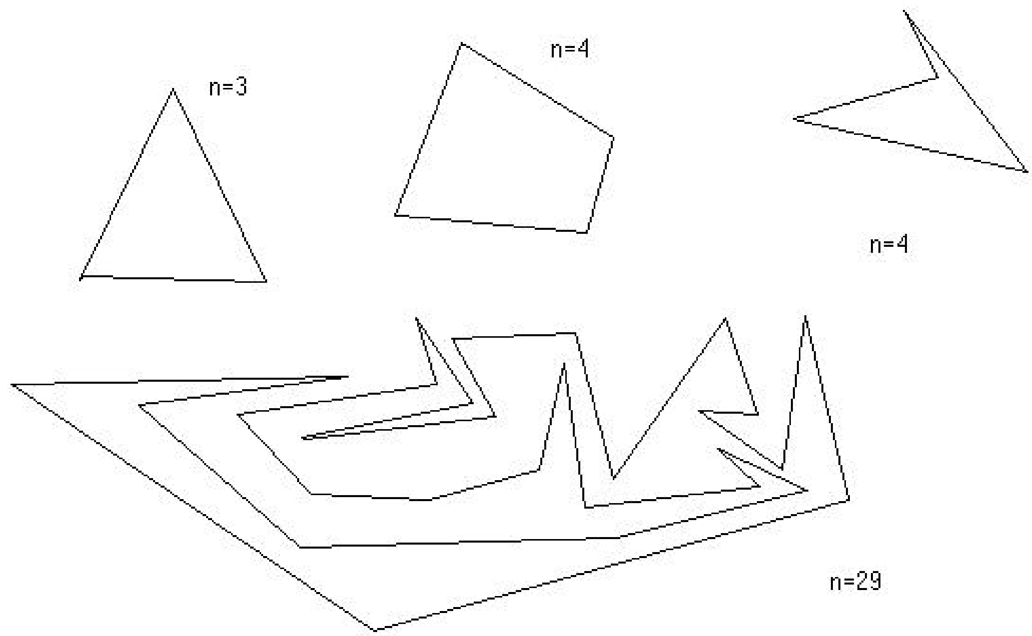

Figure 3: There is one 3-gon, 2 4-gons and who knows how many 29-gons. One 29-gon is shown above.

It is worthwhile looking at shapes of multiply connected sets. Self-intersecting closed curves enclose multiply connected regions. Such a general polygon with self-intersecting boundaries may be visualized as a collection of several simply connected n-gons, glued together at self-intersecting points like the one in Figure 4, where x is one of the ten intersection points of two edges of the polygon. The interior of the polygon in Figure 4 has eleven connected components; the interior cannot be shrunk to a point.

In Figure 2 we have a polygon with a single hole; this polygon is the union U of the five convex polygons. The union U has an outer boundary defined by parts of the boundaries of A, B C, D and E, and, an inner boundary bounding the hole. U is also multiply connected. Its interior, having a hole, cannot be shrunk to a point. Note that the hole may be non-convex as we see in Figure 2.

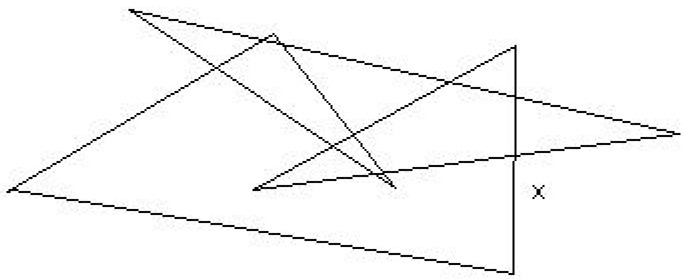

Figure 4: This self-intersecting general polygon has a self-intersecting boundary defining 11 n-gons, of which 7 are 3-gons (two very small), 2 are 5-gons, one is a 4-gon and one is a 9-gon.

Determining properties of unions of convex objects and computing unions are some of the hardest problems in "discrete and computational geometry". Unions are however of central interest in many applications. Unions of triangles in space can represent any three-dimensional structure with plane faces and straight edges. Unions of convex polygons (polyhedra) in the plane (space) can represent any two-dimensional (three-dimensional) object. Computing decompositions of non-convex objects into convex pieces is NP-hard in most cases; this happens if you wish to minimize the number of convex pieces into which you wish to partition or decompose the non-convex object. Computing unions of a given set of convex polygons or objects may not be so hard after all, not NP-hard by any means. Then what is so hard about computing unions? Well, the question is very simple: can you compute unions optimally? If the union has a certain "output size" in terms of numbers of edges and vertices, then can you compute the union in time proportional to the number of edges and vertices in the output union polygon(s)? In other words, you wish to compute the union in time proportional to the amount of "data structure" space required to store the union. You may be given a set of k convex polygons with a total of n vertices (over all the k convex polygons) as input. The k polygons may be having very complex intersections between themselves but one of them may contain all the rest of the k-1 convex polygons in its interior, in which case, the union is simply that containing polygon. The algorithm might spend a lot of time computing the union by considering the interaction between all the k polygons in the input but the answer is so simple. (Imagine for instance that a rectangle F encloses all the five convex polygons in Figure 2. Computation of the intricate union of the five polygons is not necessary to get the union of all the five polygons with F; here the union is trivially the enclosing rectangle F.) The adversary can fool the algorithm to perform more work unless the algorithm applies strategies to vanquish the evil design in the k convex polygons in the input. The breakthroughs in the 80's and 90's in data structuring, algorithms' analysis techniques, combinatorial results in geometry, and last but not the least, randomization of algorithms helped to a great extent in solving some such union problems "efficiently".

The alert and anticipating reader would not have difficulty in observing that geometric union problems arise in computer graphics and pattern recognition in many forms. To the robotics enthusiast, one main concern is how to account for obstacles and determine whether a robot can squeeze through obstacles. If the robot is not infinitesimal and in fact comparable in dimensions to those of the obstacles, then what one normally requires is computing “grown” obstacles using Minkowski’s differences, thereby reducing the robot to a point. For the case of a non-rotating robot, this is quite sufficient. Although obstacles by definition do not overlap, their grown counterparts may very well overlap, squeezing allowances and thereby preventing robot motion in certain regions. Computing unions of grown obstacles is very important indeed, as important as computing surface areas and volumes of microscopic macromolecules. Algorithms from combinatorial geometry are the best and fastest solutions for such computational problems in chemistry where scientists use various models for representing extents “volumes” of atoms and molecules or their electron clouds. Complex dynamics of large organic molecules or inorganic molecules and their rich geometries can keep computational geometers and numerical analysts busy for several years if not decades. These studies have direct consequences in biotechnology and biochemistry.

So much for the applications: we have however not yet seen much about generalizations of the very rigid definition of convexity. Can we relax convexity and still get good shapes? We did so indeed earlier in generalizing polygons from convex polygons to simple polygons as we have the 29-gon in Figure 3. What was preserved was “connectedness”; simple polygons are as “simply connected” as convex polygons are. The interiors of all such polygons can be shrunk to single point. Now we introduce another kind of generalization of convexity by using what is known as “visibility”. Visibility is a fundamental geometric concept, almost as much as convexity or connectedness is. Two points p and q are said to be (mutually) “visible” if the join pq does not intersect any object or obstacle. So, inside a convex object, every point is visible from every other point. There is total visibility in convex objects. In simple polygons, some pairs of vertices are not mutually visible because their respective joining segments can intersect the polygonal boundary. In the 29-gon in Figure 3, the leftmost and the rightmost vertices are not mutually visible, whereas, the bottommost vertex is visible from the leftmost and rightmost vertices as well as from exactly four more vertices (which ones?). As per this definition of visibility, we can say equivalently that a convex polygon is “illuminable” or visible or “seen” from any and every point inside the polygon. A simple polygon P too may be visible or illuminable or “guardable” from some interior point p as we see in Figure 5. Some polygons are not illuminable in this sense but are illuminable by a tubular source of light pq (see polygon in Figure 6).

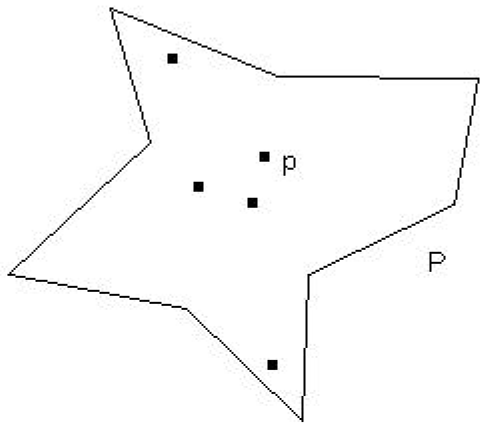

Figure 5: The topmost and bottommost interior points can not illuminate P. All the three others inluding p can see or illuminate P.

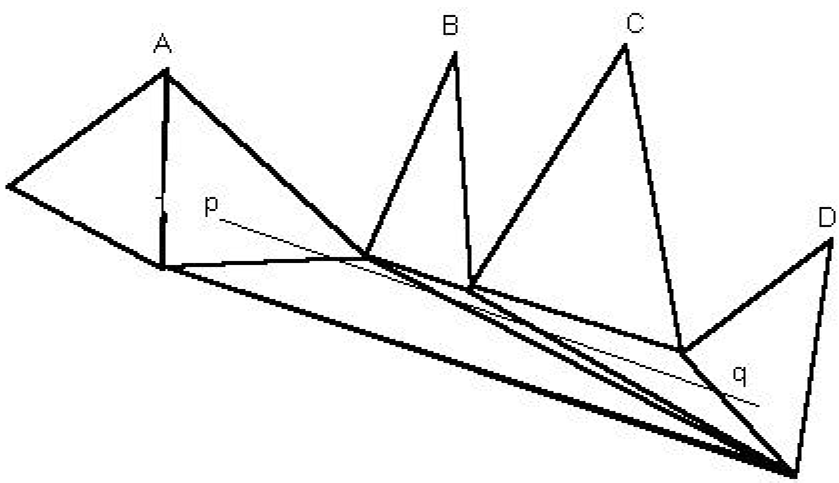

Figure 6: Guards at A, B, C and C suffice to see each point inside Q. However, even just three vertex guards suffice!

The polygon in Figure 6 is illuminable using "guards" stationed at some three vertices. Which three? Guarding or illuminating a polygon with point light sources restricted to the vertices is an academic exercise of some significance and simplicity. The polygon in Figure 6 has been purposely triangulated using edges connecting vertex pairs so that no triangulation edge intersects another. Being n-sided, there can be exactly n-2 triangles in any of the exponentially many possible triangulations of an n-sided simple polygon. So, there are exactly n-3 triangulation edges. You could triangulate the same polygon in so many ways. For instance, instead of connecting A to a lower vertex, you could have “flipped” that edge with an edge connecting the two boundary neighbors of A, again getting a valid triangulation. It is not tough to show that n/3 vertex guards suffice to guard an n-sided simple polygon. You may not need so many guards for every n-gon and computing the minimum number of vertex guards required to guard a given n-gon is an NP-hard problem. Is the corresponding decision version NP-complete? Yes, indeed. The proofs are not very hard, just some geometric constructions!

Even though an n-gon may require n/3 vertex guards, just one mobile guard or tubular light source may be sufficient for “seeing” or guarding the whole interior as we observe in Figure 6. The segment pq can see the whole polygon. Can you imagine or construct simple n-gons where one may require a large number of mobile guards, asymptotically increasing with the number n of vertices? Is the problem of computing the minimum number of mobile guards necessary for a given n-gon (the mobile guards may patrol in the interior of the n-gon), a problem in NP? Is it also in P or is it NP-complete? Well, this is not quite an easy question!

Can convexity help in guarding? For instance, is it easier to find the minimum number of vertex guards in an n-gon if it is “weakly visible” from an internal segment? Here, we use the terminology “weakly visible” from a segment to mean that the n-gon can be guarded by a mobile guard moving along the stretch of the segment. So, our question is whether computing the minimum number of vertex guards is in P (instead of being NP-hard), for n-gons that can be guarded by a single internal mobile guard. The answer unfortunately is in the negative!

Convexity may not help in “guarding” but it does certainly help in “hiding” people in n-gons. Consider the NP-hard problem of determining the maximum cardinality set of vertices of an n-gon that are not mutually visible, pair wise. A set of vertices where each pair is mutually invisible is called a “hidden set”. This is somewhat like the independent set problem on graphs. If the n-gon has certain “good convexity” properties then it succumbs to a dynamic programming approach, yielding a polynomial time algorithm for computing the maximum hidden set. The good convexity property we are alluding to is grossly the property that the (internal) shortest path between any pair of vertices in the polygon is “convex”, either always right turning or always left turning. See Figure 7 and get a feeling that the shortest path between any pair of vertices is indeed “convex”. Simple polygons that satisfy this convexity property of shortest paths can be “characterized” as simple polygons where the entire polygon is guardable by a “mobile guard” patrolling on an edge of the polygon. In other words, such polygons are weakly visible from an edge of the polygon, edge vw in Figure 7.

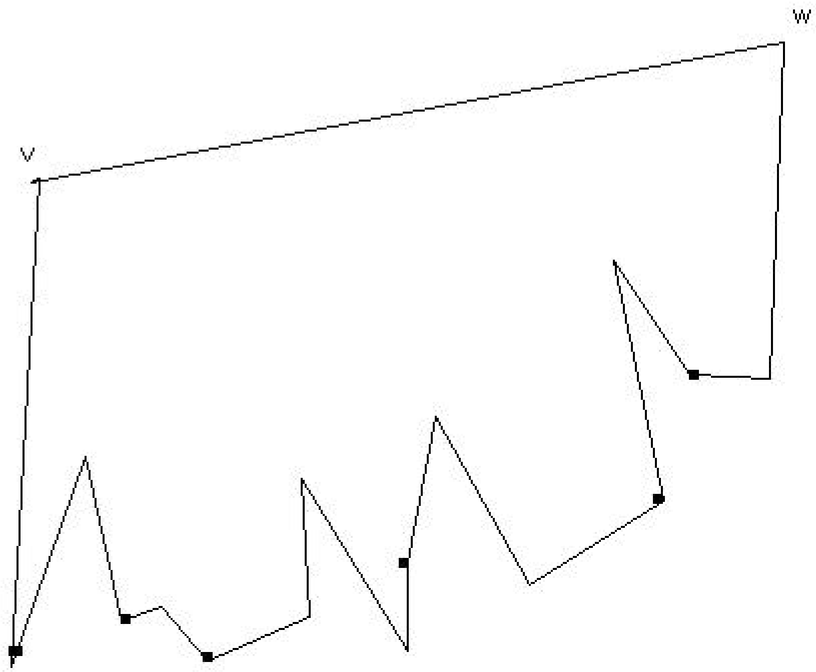

Figure 7: This polygon is weakly visible from vw. In other words, the polygon can be guarded by one mobile guard on the boundary edge vw. The bold marked vertices form a large hidden set. Is there a larger hidden set?

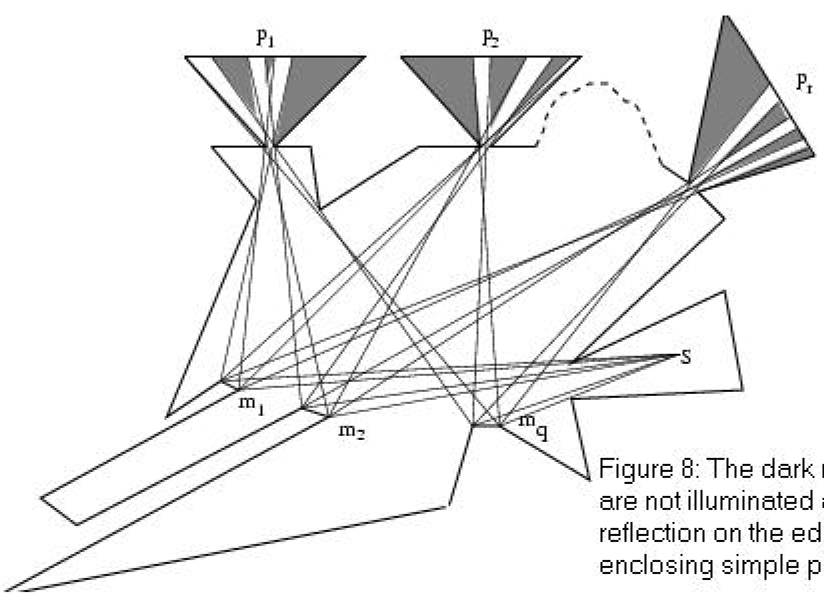

Figure 8: The dark regions are not illuminated after one reflection on the edges of the enclosing simple polygon.

Weakly visible polygons interact interestingly inside simple polygons. The simple connectedness of the enclosing n-gon imparts cute properties to weakly visible polygons generated inside the n-gon. See for instance the simple polygon in Figure 8 with q illuminated edges m1, m2, … , mq. Notice that most of the enclosing simple polygon P (an n-gon), is flooded due to light reflected ("diffuse" reflection) from various edges. The only source of light is a point source stationed at the interior point s. The point s directly illuminates only a small portion, certainly not the triangular pinhole components p1, p2, …, pr. See how carefully s is hidden from the pinhole components. Edges m1, m2, …, mq however are flooded with light from the point source at s and these edges in turn, by virtue of diffuse reflection, illuminate portions of the pinhole components and other regions of the enclosing polygon P. The walls (edges) of the polygon P are opaque but reflecting. (For a change, the reflections here are not Newtonian but Lambertian or "diffuse"; this means (as you have already rightly guessed), the reflecting edges scatter back incident light in all possible directions). The construction is cleverly done so only edges m1, m2, …, mq and no other edges illuminate the pinhole components p1, p2, …, pr. Each of the q edges send light into each of the r pinhole components, creating alternately illuminated and dark regions. Note that we entertain only one reflection after light emanates from s. Also note in that other edges appearing between in the chain of the q reflecting edges m1, m2, …, mq, are so tilted that light reflected back by them miss all the r pinholes. In summary, we get r*q=O(n^2) dark regions and thus the boundary of the region V1(s), visible from s after one diffuse reflection, has O(n^2) edges. We could chose the enclosing n-gon P to be such that q=n/c and r=n/d for constants c and d independent of n. So, for asymptotic n, the number of edges on the boundary of the region V1(s), visible from s after one diffuse reflection is O(n^2). This is interesting because V1(s) has quadratic combinatorial complexity, whereas, the region V(s), directly visible from s, can have at most O(n) edges. The linearity of V(s) is not hard to prove. It is also not hard to show that V(s) is simply connected. (Note that V(s) and V1(s) are by definition contained inside the enclosing n-gon P because P has opaque but internally diffusely reflecting edges.) However, the proof that V1(s) is also simply connected is not surprising but certainly non-trivial! After all, diffuse reflection floods light all around points of reflection and therefore invisible regions are tough to create, particularly those dark regions that are totally surrounded by illuminated regions.

You might have been wondering how there can be holes inside visible regions. We have not seen any so far. V(s) and V1(s) above do not have holes. The dark regions in Figure 8 are not holes; those dark regions are peripheral and if you land up inside one of the O(n^2) dark regions of P, you can exit P without having to pass through any point of V1(s). So, you can escape radiation from s even if the radiation is after one reflection! You are simply not trapped. Having a hole inside a visible region, as mentioned in the previous paragraph is interesting because once you are inside such a hole or island of darkness, escaping to the exterior of P requires you to drive through some portion of the visible region, irrespective of which way you take a continuous drive out of the hole to the exterior of P. This would have happened for holes of V1(s), if there were any. Fortunately, this is not possible. What about V2(s)? We define V2(s) to be the sub-region of P that is visible from s after at most two diffuse reflections on the wall of P. All walls (edges) of P are diffuse reflectors. So, for any point q in V2(s), there is a polygonal path from s to q suffering at most two reflections on edges of P. (Remember that every point p of P inside V1(s) is such that there is a polygonal path from s to p with at most one reflection on an edge of P as you see in Figure 8.) It turns out that V2(s) is not simply connected, has holes, and quite plenty of them and therefore we say that V2(s) is multiply connected. This is a bit surprising, indeed true, but not quite easy to show. We omit our discussion on diffuse reflections here and get to Newtonian reflections below to show you holes in visible regions.

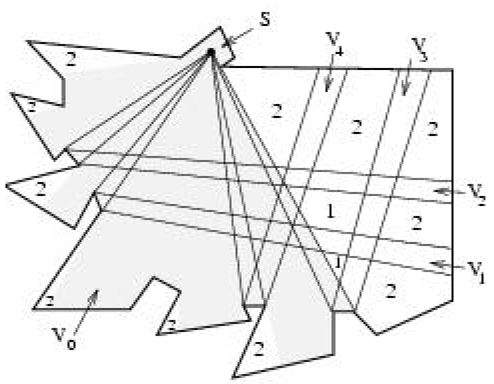

Figure 9: Specular or Newtonian reflections create horizontal beams V1, V2, and vertical beams V3 and V4, creating a dark hole labelled 1 and peripheral dark regions labelled 2. See how nicely the hole is surrounded by regions visible from s after one specular reflection.

In Figure 9, reflections are specular; the angle of incidence is equal to the angle of reflection. This is different from the case of diffuse reflections. (Specular reflections happen on mirrors you use to see your face. Diffuse reflections happen anywhere and everywhere you have partly rough surfaces like walls and furniture.) The point s in Figure 9 is again nicely hidden from the region having the dark regions marked 1 and 2. The region hidden from s gets light after one reflection only across beams V1, V2, V3 and V4. Other edges reflect light coming from the only point source s but the rays so resulting have to suffer one more reflection to reach anywhere; as we are interested in only one reflection now, we ignore such rays. If you had n/c beams like V1 and V2 and n/d beams like V3 and V4, then you would have got O(n^2) holes like the region marked 1. So, we can now see plenty of holes in the region visible from s after at most one specular reflection. We say that V1(s) (for specular reflections) is multiply connected, as was V2(s) in the case of diffuse reflections. It turns out that the number of holes increases exponentially (in the case of specular reflections) with the number k of permitted reflections. We have so far been looking at union problems, unions of convex polygons and unions of regions visible from edges and points, as in the last few paragraphs. The richness in these problems is due to the combination of convexity, connectedness and visibility.

We now try to take a look at some elegant intersection properties of convex sets. Suppose n people attend a party as per the following rule: each person enters the party exactly once and leaves the party exactly once. To determine whether all the n persons were simultaneously at the party at some moment of time, it suffices to show that each person's party visit time interval overlapped with each other person's party visit time interval. In other words, we have an O(n^2) time check and do not need to actually compute the complex intersection of all the n time intervals to check its non-emptiness. (It is quite easy though to compute the actual intersection of all the n time intervals in O(n log n) time!) Now get down to two dimensions: take three rectangles A, B and C and place them on your table. Try placing them so A intersects B, B intersects C and C intersects A. Do all of A, B and C have a common intersection? May not be, unless you keep the three rectangles oriented in such a way that their edges are parallel to the X- and Y- axes. Suppose you had n convex objects in the plane. These could include triangles, circles, rectangles or any kind of convex objects. Suppose you place them on the plane. How do you know that all the n intersect at a common point? As we have already seen, checking pair wise does not help. Can we check all triplets? Try verifying the following result: if every triplet of the given n convex objects have a common intersection, then all the n objects have a common intersection. Believe it or not, this result is very easy to prove as much as its d–dimensional counterpart is: given n>d+1 convex objects in d-dimensional space, the n objects have a common intersection if each of the (d+1)-sized subsets of the n objects have common intersections. So, you can see how degenerate matters become as we asymptotically increase the number n of objects, keeping the dimension fixed at d. However, note that this result is very fundamental and significant indeed with very far reaching implications. For instance, this result helps in one way of proving the famous and heavily used Lipton-Tarjan separator theorem for planar graphs.

Another consequence of the above result is that in order to determine whether a set of n points in the plane can be covered by some unit disc, all you need to do is check that each triplet of n points can be covered by some unit disc. So, an O(n^3) algorithm can decide if n points in the plane can be covered by a unit disc. Try showing that checking each pair of points does not help.

Intersection and union problems are fundamental to geometry and the study of their combinatorial and computational aspects give us insights into how the whole is made out of its parts. Convexity, connectedness and visibility are important geometric features that guide our way through the study of shapes.

-------------------------------------- The End ---