GALLERIES AND LIGHT MATCHINGS: FAT COOPERATIVE GUARDS

by

Jurek Czyzowicz
Département d'Informatique
Université du Québec à Hull
Hull, Québec

and

Ivan Rival and Jorge Urrutia
Department of Computer Science
University of Ottawa
Ottawa, Ontario
Canada

Abstract. For any collection of n disjoint line segments on the plane, such that no two are parallel and no three extensions meet in a common point, $\lfloor 3n/4 \rfloor$ lights, none on any of the line segments, are occasionally necessary, and always sufficient to illuminate all points on all of the line segments.

AMS subject classifications (1980) 05B25, 51E30, 05C15.

Key words and phrases: Line segments, convex partition, illumination, planar graph, colouring.

GALLERIES AND LIGHT MATCHINGS
I. FAT COOPERATIVE GUARDS

by
Jurek Czyzowicz, Ivan Rival and Jorge Urrutia

How many guards are necessary, and how many are sufficient, to patrol the n paintings of an art gallery—especially a modern one, with its numerous alcoves, corners, and narrow snake-like passages? Fuelled by the modern interest in computational geometry, this wonderfully naïve question of combinatorial geometry has, since its formulation [V. Klee (1973), cf. R. Honsberger (1976)], stimulated a rush of papers, surveys, and even a book [J. O'Rourke (1987)], most written in the last decade. The mathematical beginnings are found in the well-known gem of a theorem by V. Chvátal (1975) according to which $\lfloor n/3 \rfloor$ stationary guards are occasionally necessary and always sufficient. A *guard* or simply, a *light* is a stationary light source which can survey 360˚ about its fixed, designated position. We shall suppose that guards must be placed well away from line segments. To this end we suppose that a light is a disk in the plane with a non-negligible radius whose light source lies at its centre. Thus if one endpoint of a line segment is closer than this radius to an interior point of another line segment then no guard can fit in that gap between the two line segments. Still, we shall suppose that a light's radius is nevertheless sufficiently small that it at least fits on either side of a point on a segment not itself on the extension of another line. In this formulation the gallery, having straight walls, is a polygon of n vertices and, once illuminated, there is, for every point of the polygon and its interior, a light (guard) which illuminates it.

There are numerous other interesting variations: for instance, traditional orthogonal art galleries all of whose walls are either horizontal or vertical, in which case $\lfloor n/4 \rfloor$ guards are necessary and sufficient [J. Kahn, M. Klawe, and D. Kleitman (1983)]; mobile guards, each of whom patrols from along a line segment within an n-vertex polygon, of whom $\lfloor n/4 \rfloor$ are necessary and sufficient [J. O'Rourke (1983)].

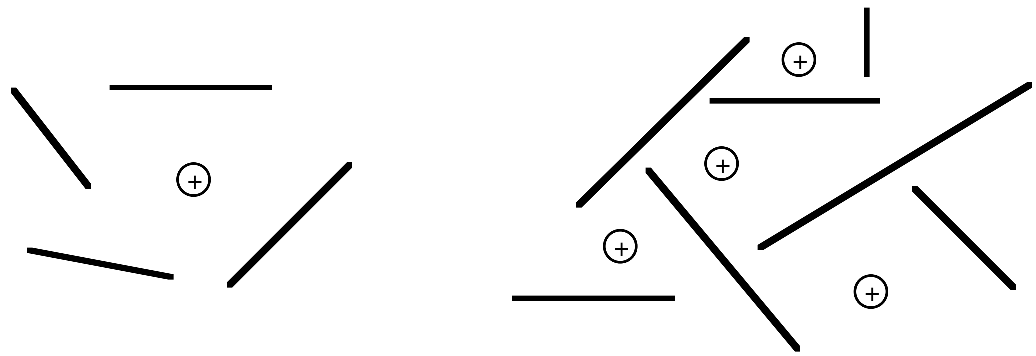

One light sees all. Four lights see all points.

Figure 1

We consider a collection of n disjoint line segments on the plane, which need not, however, form a polygon. In fact, rather than a wall of the art gallery, we may consider any line segment to represent the actual *objet d'art*, a painting or sculpture, say. And, rather than confining the gallery within a polygonal region, we take it to be the entire plane, thus, a spacious "open-concept" gallery. Given such a collection of line segments we say that a light placed at point x *sees* or *illuminates* a point y on the plane if the line xy does not cross the interior of any other line segment.

How many such lights are needed to illuminate all of the points on all of the line segments? Here is our main result.

THEOREM 1. *For any collection of* n *disjoint line segments on the plane, such that no two are parallel and no three extensions meet in a common point,* $\lfloor 3n/4 \rfloor$ *lights, none on any of the line segments, are occasionally necessary, and always sufficient to illuminate all points on all of the line segments*.

The parallelism and intersection constraints on the family of line segments do not seem to be serious in the physical model for, a little "shaking" of any collection of n disjoint line segments will certainly misalign any unwanted symmetry.

There is an apparently related result of J. O'Rourke (1987) according to which $\lfloor 2n/3 \rfloor$ lights are necessary and sufficient to illuminate all points on the plane. In that result, lights are occasionally placed right on points belonging to the collection of line segments so such lights can thereby see points of the plane on either side of the line. In that model the n line segments are obstacles to the illumination of the plane. Our own model differs in several respects. The lights are not intended to illuminate all points on the plane—just the points on the line segments. Thus, our guards patrol the paintings and sculptures without necessarily illuminating any seat, lounge, or wallflower. In addition, none of our lights may be placed on a line segment—paintings must remain uncluttered.

Elsewhere [S. Foldes, I. Rival and J. Urrutia (1988)] we have considered the order-theoretical structure of the "blocking relation" induced by a single light source which, roughly speaking, radiates through the successive layers of the line segments.

Here is an algorithmic companion of Theorem 1.

THEOREM 2. *Given any collection of* n *disjoint line segments on the plane, such that no two are parallel and no three extensions meet in a common point, there is a polynomial time algorithm assigning* $\lfloor 3n/4 \rfloor$ *lights, none on any of the line segments, which illuminate all points on all of the line segments. Moreover, there is a data structure to support any such configuration of* n *disjoint line segments which can be constructed in time* $O(n \log n)$ *and there is an* $O(n)$ *time algorithm to assign* $\lfloor 4n/5 \rfloor$ *lights to illuminate all points on all of the line segments.*

It is natural, too, to ask about this analogous problem: how many lights are needed to illuminate both sides of the n lines? In this model, each point needs to be illuminated twice, once on each side of its line. In fact, once an appropriate, associated data structure is generated this problem is easy to solve.

THEOREM 3. *Given any collection of* n *disjoint line segments on the plane, such that no two are parallel and no three extensions meet in a common point,* $n+1$ *lights, none on any of the line segments, are occasionally necessary and always sufficient to illuminate all points on both sides of all of the line segments. Moreover, this assignment can be carried out in* $O(n \log n)$ *time.*

The problems we have considered here involve two assumptions. In the first place, we have assumed that the guards are "cooperative", in the sense that the illumination of a painting may be shared by several of them. It is a much harder problem to minimize the number of lights located at points on the plane so that, for each line, there is at least one light which fully illuminates it. For a point x on the plane, say that x *sees* or *illuminates a line segment* S if x illuminates each point y on S. In the second place, we have assumed that the lights be placed well away from the line segments, for the light source itself is encircled by a disk of (constant) nonzero radius. Thus, our guards are "fat". In particular, as long as no endpoint is too near an interior point of another segment, a light may placed near the point of intersection of the extension of the second with the first. Such a light will illuminate all points on both of these lines. If, on the other hand, the lights are indeed point sources not requiring any nontrivial disk, that is, "thin" guards, then they may be placed even at the apparent junction of two line segments which, when extended, intersect. The case of "thin responsible" guards is different. Elsewhere [J. Czyzowicz, I. Rival and J. Urrutia] we have shown, for example, that *for any collection of* n *disjoint line segments on the plane, such that no two are parallel and no three extensions meet in a common point,* $\lfloor n/2 \rfloor$ *lights are occasionally necessary and* $\lfloor 3n/4 \rfloor$ *lights are always sufficient to illuminate all of the* n *line segments.*

Proof of Theorem 1. One proof, due to S. Fisk (1978), of the original $\lfloor n/3 \rfloor$ art gallery theorem consists, in outline, of these three steps:

1. triangulate the polygon;
2. three-colour the corresponding graph;
3. place a guard at each vertex labelled by the least popular of the three colours.

Our proof of Theorem 1 sets out with an analogue of triangulation. Extend each line segment, one at a time in both directions until it hits either another segment or a previous segment extension. This yields a partition of the plane into $n+1$ convex regions, a simple fact, noted before in the literature, but whose proof requires a measure of subtlety. For instance, extend one line segment as far as possible in both directions (infinite in at least one direction). Then extend every other line segment until it just touches another. Once done, the configuration of line segments so prepared is "connected". Then, each further extension of any line segment to hit another line segment will always split a region into two. Of course, the actual induced convex partition depends on the order in which the extensions are drawn (see Figure 2).

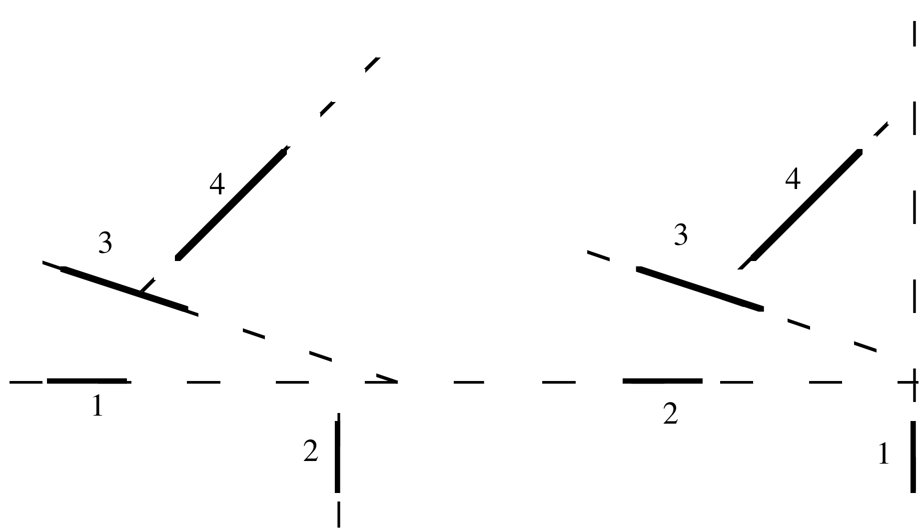

Figure 2

Next construct the dual graph of the graph corresponding to the convex partition, that is, assign a vertex to each of the $n+1$ convex regions and join two by an edge if their regions share a common boundary. As no three of the line extensions meet in a common point, this dual graph is a (triangulated) planar graph. In particular, its vertices can be four-coloured. The largest colour class has at least $\lceil (n+1)/4 \rceil$ vertices. Finally, assign a guard to all *other* vertices, at most $\lceil 3n/4 \rceil$ of them.

If there is a point on a line segment not illuminated then it must be contained in a convex region of the partition which is not assigned a guard. As each adjacent region is, however, assigned a guard it does illuminate the common boundary. Thus, all points on all line segments are duly illuminated.

Suppose, on the other hand, that the points of the line segments are already illuminated by a minimal distribution of lights on the plane. Each such light corresponds to a vertex of the dual graph and, moreover, each edge of the dual graph is incident with one such vertex. Thus, these distinguished vertices form an *edge cover* of the dual graph. The dual graph of the collection illustrated in Figure 3 is the complete, four-vertex graph K_4. Any edge cover of it requires three vertices $(3 = \lceil 3 \cdot 3/4 \rceil)$.

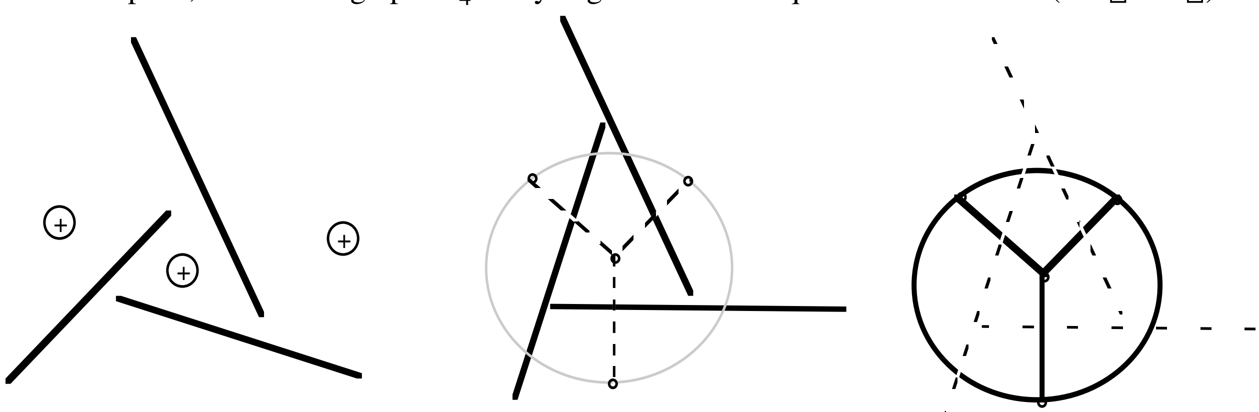

Three lights suffice | An induced convex partition | The dual graph K_4

Figure 3

This idea we can use to construct a graph on n vertices, which is the dual graph of a convex partition induced by such a corresponding collection of n disjoint line segments and which graph itself has minimum edge cover of size $\lceil 3n/4 \rceil$

For instance, if $n = 4m$ we can construct such a graph on n vertices, consisting of m disjoint copies of K_4, arranged on the plane and with extra edges added to form a triangulation. This graph has minimum edge cover of size $3m$ (for each K_4 itself must have an edge cover). (See Figure 4 for an example.) Now the dual of this graph is a planar embedding of a graph (not necessarily with all edges straight) which, if suitably perturbed, has only straight edges (cf. Fary (1948)).

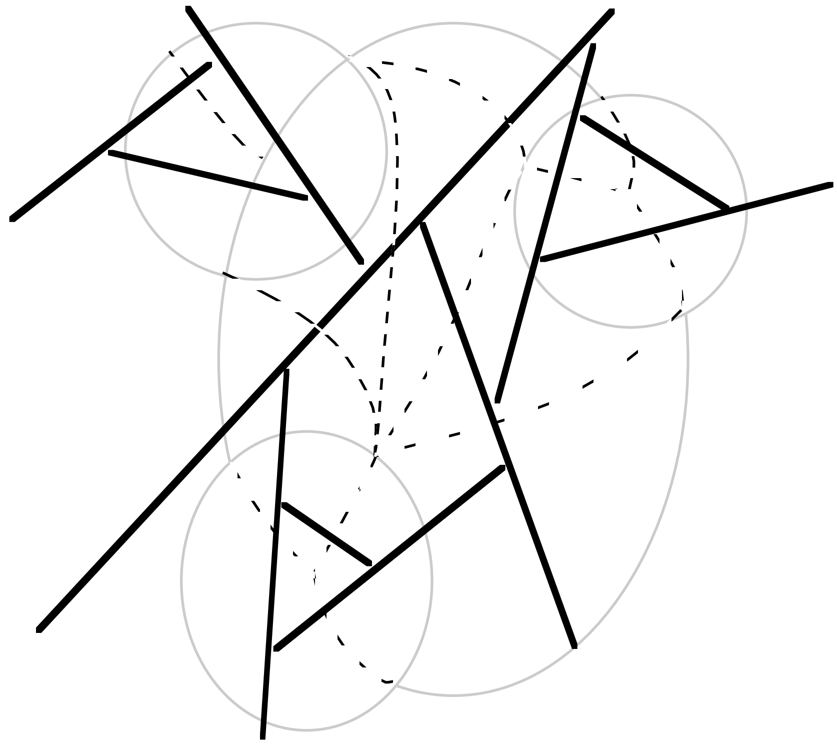

n = 12

Figure 4

That is, this dual graph is a convex partition of the plane induced by a collection of n disjoint line segments. If $n = 4m + r$, $0 < r < 4$, first use a triangulation of m copies of K_4, then adjoin r more vertices as a K_r, and triangulate the result. Its dual is a graph whose edges, once "straightened", forms a convex partition of the plane.

Proof of Theorem 2. The first step in the assignment of lights is the construction of a convex partition of the plane from the initial collection of n disjoint line segments. We may suppose that none of them is vertical. (The analogous step for the $\lfloor n/3 \rfloor$ art gallery theorem entails a triangulation on an n-vertex polygon, for which the best known algorithm requires $O(n \log \log n)$ time [R.E. Tarjan and C.J. van Wyk (1986)].) To implement a convex partition of the plane we may apply the "sweeping approach" according to which a vertical line, the *sweep line*, gradually moves from left to right across the plane recording, at each position, all information about the problem to the left of the sweep line relevant to solving the problem to the right of the sweep line. It uses a data structure, called the *y-structure*, to record the status of the sweep at the current position of the sweep line. The algorithm involves extending some of these line segments, first to the right (and, in a subsequent iteration, to the left). Any point on an extended line is either *active* or *dead*. Every point on any of the initial line segments is active. The line segments intersecting the sweep line in active points are sorted by their y-coordinates. The sweeping approach also uses another data structure, the x-*structure* to record *real* and *potential* x-coordinates. The x-coordinates of the endpoints of all initial line segments are real and are sorted along this x-structure. As the sweep line moves from left to right, only a few of its positions can actually cause a change of its status—reaching an x-coordinate of a real point which, in this first iteration consists entirely of the x-coordinates of endpoints of initial line segments.

Given a position of the sweep line, consider (with respect to the y-structure) any consecutive active points whose corresponding line segments L, K intersect to the right. Extend both of the line segments L, K to the right. If their common point is an interior point of the initial segment L, say, then call all points in the extension of K beyond it dead. (Of course, the intersection point cannot belong to both of the initial segments.) If the common point is not interior to either L or K then call this common point an active point, of L, say, in which case again, all points on the extension of K beyond it are dead. The common point of L and K is a potential point. Once done for every consecutive pair identified by the sweep line, move the sweep line to the right until the status of its y-structure changes for the first time again.

Once a full sweep from left to right is completed, label all potential points corresponding to active points of lines as real points. Update the x-structure by adjoining these new real points and sort them all.

There may still be line segments opening to the right, whose extensions, therefore, intersect to the left. We may now repeat the sweeping approach, this time, from right to left.

The convex partition consists of all such constructed lines consisting entirely of active points only. Each step of the sweeping approach involves the construction of intersection points for consecutive pairs of lines (already sorted along the y-structure), whose time required is $O(\log n)$. The most expensive of the operations used in the sweeping approach is the sorting of the x-structure and, even with the updating, uses only time $O(n \log n)$. The entire iteration is repeated from right to left requiring in total time $O(n \log n)$. (For further details of this approach see [Mehlhorn (1984)] or [Preparata and Shamos (1985)].)

The second step in the algorithm depends on the complexity of the colouring algorithm used. A four-colouring is involved for the □3n/4□ estimate and, while polynomial in time complexity, no truly efficient general four-colouring scheme is currently known (cf. [Appel and Haken (1977)]). On the other hand, a five-colouring of a planar graph can be carried out in linear time ([Chiba, Nishizeki and Saito (1980)]).

Proof of Theorem 3. Construct a convex partition of the plane from the n disjoint line segments. (This, too, is an instance of the problem.) This partition is constructed in $O(n \log n)$ time, as we have just seen in Theorem 2. It produces $n + 1$ regions on the plane, some bounded, others unbounded.

For each of these regions a light must be placed inside it for the inner sides of the bounding lines to be illuminated and, in fact, one light inside illuminates all of the interior sides, for each region is convex.

REFERENCES

K. Appel and W. Haken (1977) Every planar map is four-colourable, *Illinois J. Maths.* **21**, 429-567.

N. Chiba, T. Nishizeki and N. Saito (1980) A linear algorithm for five-colouring a planar graph, in *Lecture Notes in Computer Science* **108**, *Graph Theory and Algorithms*, Springer-Verlag, pp. 9-19.

V. Chvátal (1975) A combinatorial theorem in plane geometry, *J. Comb. Theory, Ser. B* **18**, 39-41.

J. Czyzowicz, I. Rival and J. Urrutia (1989) Galleries and Light Matchings II. Thin Responsible Guards, preprint.

I. Fary (1948) On straight line representation of planar graphs, *Acta. Sci. Math. Szeged* **11**, 229-233.

S. Fisk (1978) A short proof of Chvátal's watchman theorem, *J. Comb. Theory, Ser. B* **24**, 374.

S. Foldes, I. Rival and J. Urrutia (1988) Light sources, obstructions and spherical orders, University of Ottawa Technical Report TR-88-23.

R. Honsberger (1986) *Mathematical Gems II, Math. Assoc. of America,* 104-110.

J. Kahn, M. Klawe and D. Kleitman (1973) Traditional galleries require fewer watchmen, *SIAM J. Alg. Disc. Math.* **4**, 194-206.

K. Mehlhorn (1984) *Multi-dimensional Searching and Computational Geometry*, EATCS Monograph on Theoretical Computer Science, Springer-Verlag.

J. O'Rourke (1987) *Art Gallery Theorems and Algorithms*, Oxford University Press.

F.P. Preparata and M.I. Shamos (1985) *Computational Geometry, An Introduction*, Springer-Verlag.

R.E. Tarjan and C.J. Van Wyk (1986) An O(n log log n)-time algorithm for triangulating simple polygons, ITT Bell Laboratories Technical Report.