The Minimum Cone-Segment Cover Problem

Naji Mouawad
Department of Computer Science
University of Waterloo
Waterloo ONT. N2L 3G1
Email: nmouawad@watmath.waterloo.edu

Abstract

Given a source of X-rays emitting in a 'conic' fashion, one would like to construct a safety barrage that could block any accidental high emittance of radiation. The barrage consists of panels coated with lead to be placed within the cone of emission. Due to space constraint these panels can be put only along specific lines. If one wishes to minimize the total surface of the panels, then its argued that one panel is not always the best solution.

We formalize the problem as follows: let E^2 be the Euclidean plane, O its origin. Let C be a cone of E^2 with apex O, and bounding rays r and r'. A set of segments S *subtend* C if the endpoints of the segments of S lie on r and r'. Let L be a line intersecting r and r' such that O and S belong to the bounded subset of E^2 determined by L, r, and r'. Let A be the unbounded subset of E^2 also determined by L, r, and r'. A minimum length-cover of O given C and S, is a set of sub-segments of S with minimum length which hides O from A.

We exhibit an optimal $\Theta(n \log n)$ algorithm for finding such a minimum cover when the edges of E are non-intersecting, and an $O(m \log m + mn \log n)$ algorithm when the number of intersection points among the edges of S is m.

1 Definitions

Let h_0 and h_1 be two closed half-planes whose respective bounding lines l_0 and l_1 intersect in O. The boundary of the intersection region I of h_0 and h_1 consists of two closed rays r_0 and r_1, such that r_0 lies on l_0 and r_1 lies on l_1, and r_0 and r_1 share the common endpoint O. With respect to O, the ray of $\{r_0, r_1\}$ bounding I, in the clockwise direction from I shall be called the *right* bounding ray. The other ray shall be called the *left* bounding ray. Without loss of generality, we shall assume that r_0 is the left ray and r_1 the right ray. We shall refer to the region I as the *cone* $K = (r_0, r_1)$ determined by r_0 and r_1. In general, any two rays sharing the origin as their common endpoint and not contained in a common line define a unique cone.

Given a cone $C = (r, r')$, the *angle* of C is defined to be the angle $\angle rOr'$, formed by r, O, and r'. Observe that this angle must always be less than π. The *apex* of any cone is the origin O. Given two cones C and C', if $C' \subseteq C$ then C' shall be said to be a *sub-cone* of C, Two cones C and C' will be considered *disjoint* if the apex O is the only point that belongs to both cones, *contiguous* if they intersect in a single ray, and *overlapping* otherwise.

Let C be a cone in E^2 with apex O and bounding rays r and r'. The bisector of the cone C is assumed to lie on the positive x-axis; all other cases are treated in a similar manner with a minor re-parameterization. Let C' be the symmetric cone of C with respect to O. Let Q be the set of points between C and C' with positive y-coordinate, Q' the set of points between C and C' with negative y-coordinate.

Let E be a set of edges subtending C. Given an edge e of E, a continuous interval of e is called a *segment* of e. Let e and e' be two subtending edges over cone C. If the segment of e subtending an arbitrary sub-cone C' of C, is shorter than the corresponding segment of e', then e *dominates* e'. If none of the edges dominates the other, they are said to be *incomparable*. In the latter case, the minimum cover determined by these two edges is composed of two segments, each of which belongs to one of the edges. The ray separating the two segments is called the *splitting ray* of e and e'.

2 The non-intersecting case

2.1 Linear Upper Bound

Let C be a cone with two subtending edges e and e' with e dominating e'. What happens if we 'expand' the cone and the edges ? Is the relationship between e and e' preserved ?

Lemma 2.1 *Let C be a cone divided into three contiguous sub-cones C_L, C_M and C_R, termed left, middle and right sub-cones. Let e_1 and e_2 be two subtending edges of C. Denote by $e_{i,J}$ the segment of edge e_i subtending sub-cone C_J. If $e_{i,M}$ dominates $e_{j,M}$, $i \neq j$, then either $e_{i,L}$ dominates $e_{j,L}$, or $e_{i,R}$ dominates $e_{j,R}$.*

PROOF *Outline of the proof*: We express the length of $e_{i,2}$ in terms of $e_{i,J}$, where $J = L, \text{or} R$; the choice of J depends on the sign of the slopes of e_1 and e_2, and the location of the intersection point of their underlying lines in any of the two quadrants Q or Q'. For instance, if both edges are positively sloped and the intersection point is in Q', and we assume that $e_{1,M}$ dominates $e_{2,M}$, then $J = R$.

Let $L(e_{i,M})$ be the length of $e_{i,M}$. We obtain the following relation :

$$L(e_{i,M}) = W_i.L(e_{i,J})$$

Where $W_i > 0$. By hypothesis $e_{i,M}$ dominates $e_{j,M}$, thus $L(e_{i,M}) < L(e_{j,M})$. Then it is sufficient to proof that $W_i > W_j$. Which is straightforward to establish. □

If both $e_{i,M}$ and $e_{i,L}$ are dominant, we say that $e_{i,M}$ is *left extensible*, similarly if $e_{i,M}$ and $e_{i,R}$ are both dominant, $e_{i,M}$ is *right extensible*.

If we know that $e_{i,L}$ and $e_{i,R}$ are dominant over their respective sub-cones, what can we say about $e_{i,M}$?

Define the *Hull* of two cones to be the smallest cone containing them.

Lemma 2.2 *Let C_1 and C_2 be two disjoint cones, H their hull. Let e_1 and e_2 two subtending edges of H, such that $e_{1,i}$ dominates $e_{2,i}$, $i = 1, 2$, then e_1 dominates e_2.*

PROOF The proof is by contradiction. It assumes that e_1 and e_2 are incomparable over H, and uses the previous lemma to obtain a contradiction by noting that if the segment of e_2 over $H \setminus (C_1 \cup C_2)$ is dominant, then it would be either left or right extensible, contradicting the dominance of $e_{1,J}$, $J = L, R$. □

Theorem 2.3 *Given a cone C and a set of n non-intersecting subtending edges, the number of segments in any minimum cover, is no more than n.*

PROOF By contradiction. Assume that there exists a minimum cone-segment cover solution M with m segments, $m > n$. By the pigeon-hole principle, there exists at least two disjoint dominant segments over cones C_j, C_k that belong to the same edge e_i. Let $e_{l,d}$ be a minimal segment between $e_{i,j}$ and $e_{i,k}$. Since both segments $e_{i,j}$ and $e_{i,k}$ are dominant, the segment of e_i subtending $hull(C_i, C_j)$ is also dominant by lemma 2.2, contradicting the fact that $e_{l,d}$ is a dominant segment. Hence each subtending edge can contribute with at most one dominant segment. This fact establishes the result. □

This upper bound is tight as established by theorem 2.6.

2.2 Outline of the Algorithm

The algorithm sorts the set E closest to furthest with respect to the origin. Next, the set E is divided into two non-intersecting subsets E_- and E_+, where E_- contains all negatively sloped edges, and E_+ all positively sloped ones. Two distinct minimum-covers, M_- and M_+, are constructed in a symmetric way, one with E_- as the set of subtending edges, the other with E_+. The minimum-cover induced by E is constructed by means of a simple merge-like technique, where segments of M_- and M_+ that cover a common area of the cone are compared against each other; the outcome of the comparison being that either one of the two segments is dominant, in which case the other one is discarded, or that they are incomparable, in which case each of the two segments contributes with one subsegment to the minimum cover M. This process is performed in one sweep with no backtracking.

Remains to show how to construct M_- and M_+. Since the construction is symmetric, we will outline the construction of M_+.

The first step removes all the edges of E_+ whose underlying line intersects with the underlying line of e_1 in Q', where e_1 is the closest edge of E_+ to O. The reason being that those edges are dominated by e_1. This can be seen by using triangular inequalities in a straightforward manner.

Let $E'_+ = \{e'_1, \ldots, e'_t\}$, be the ordered set of remaining edges of E_+. The second step constructs M_+ as follows : starting with e'_1 and e'_2, it computes their splitting ray $\rho_{1,2}$. If $\rho_{1,2}$ falls to the left of C, e'_2 is removed, if it falls to the right of C, e'_1 is removed, otherwise the segment of e'_1 to the right of $\rho_{1,2}$ and the segment of e'_2 to the left of $\rho_{1,2}$ are inserted in M_+. The same process is repeated with e'_2 and e'_3. If e'_3 is removed, we move on to e'_4, if e'_2 is removed, we move back to e'_1, otherwise, the segment of e'_2, is replaced by a two segments, one from e'_2 , the other from e'_3. At the end of this step, we obtain a minimum cover that forms a "diagonal", meaning that the segments of M_+ are monotone when viewed radially around O. Each time the algorithm backtracks while processing this step it removes one edge. Thus constructing M_+ is done in linear time.

The following two theorems establish the correctness of the previous step.

Assume that the line equation of a subtending edge e_i is given by $y = a_i x + b_i$. Furthermore,

- Let $\delta_i = b_i\sqrt{1 + a_i^2}$
- Let $\alpha = \sqrt{\left|\frac{\delta_1}{\delta_2}\right|}$.
- Let $\tau_0 = \frac{a_1 - a_2\alpha}{1-\alpha}$.
- Let $\tau_1 = \frac{a_1 + a_2\alpha}{1+\alpha}$.

Let $y = t_1 x$ be the line equation of ray r; ray r' being its symmetric with respect to the x-axis has line equation $y = -t_1 x$.

Theorem 2.4 *There exist a splitting ray if and only if τ_0 belongs to the open interval $]-t_1, t_1[$. Moreover, the line equation of the splitting ray is given by $y = \tau_0 x$.*

PROOF Let $y = tx$ be the line equation of the splitting ray, r_t. Set $C_1 = (r, r_t)$ and $C_2 = (r_t, r')$. We want to fix t such that the sum of the lengths of $e_{1,1}$ and $e_{2,2}$ is minimized, or such that the sum of the lengths of $e_{2,1}$ and $e_{1,2}$ is minimized. . Let $f(t) = L(e_{1,1}) + L(e_{2,2})$. The slope of the splitting ray is a root of f. Deriving f, and solving $f'(t) = 0$, yields τ_0 and τ_1 as roots. Further algebraic manipulations establish the result. □

Assume that the splitting ray is given by $y = \tau_i x$, $i = 1, 2$. Let $S_1 = (r_1, r_{\tau_i})$, and $S_2 = (r_{\tau_i}, r_0)$ be respectively the left and right sub-cones, determined by $y = \tau_i x$. Let p be the intersection point of the underlying lines of e_1 and e_2.

Theorem 2.5 *If p belongs to Q and there exist a splitting ray, then the pair $\{e_{1,2}, e_{2,1}\}$ forms a minimum cover.*

PROOF *Outline of the proof:* We show that $L(e_{1,2}) < L(e_{2,2})$, by simple algebraic manipulations of the expressions of $L(e_{i,j})$. □

The proof of correctness has been sketched by the various lemmas and theorems stated. The overall idea being that no edge can participate with more than one segment in the minimum solution. This result lead to the computation of the unique splitting ray between two subtending edges. By studying the relationships between the slope of the splitting rays and the slope of the edges, and the location of the intersection point of their underlying line, minimum pairs where determined in the previous theorem. The generalization of this result to n edges lead correctly to the outlined algorithm.

Aside from the initial sorting step, every other step is performed in linear time : dividing E into two sets, eliminating all the edges whose underlying lines intersect with the underlying line of e_1 in Q', constructing the sets M_+ and M_-, and finally constructing the minimum cover M.

2.3 Lower Bounds

Although the slopes of the splitting rays implies the use of the square root function, the algorithm does without it since its decisions are based on comparisons of slopes. Thus squaring the slopes of the splitting rays, a constant number of times, avoids using the square root. Moreover, the algorithm outputs a sequence of triplets of the form (e_i, e_j, e_k), each of which represents a segment in the minimum cover, determined by the splitting rays of (e_i, e_j) and (e_j, e_k). Computing these splitting rays involves some implementation of the square root function; a step that is left open to the user.

The problem of determining this sequence of triplets has an $\Omega(n \log n)$ lower bound, as established by the following two theorems.

Theorem 2.6 *Given a cone C and three non-intersecting positively sloped subtending edges e_1, e_2, e_3, with underlying lines intersecting at point p in Q, then there exist a unique minimum cover containing one segment from each edge.*

PROOF *Outline of the proof:* Assuming that the three edges are in order e_1, e_2, e_3, then by using implicit derivation, we show that the slope of the splitting ray of e_1 and e_2 is less than that of e_2 and e_3. Theorem 2.5 implies then that there are three segments in any derived minimum cover. Since the slopes of the splitting rays are fixed, there is but one minimum-cover. □

Theorem 2.7 *Finding the sequence of triplets requires in the worst case $n \log n$ time under the algebraic model of computation*

PROOF Transformation from the SUCCESSOR problem, which is defined as follows: Given a set A of n positive integers, determine the successor for each of them in A. This problem has an obvious $n \log n$ lower bound.

Map each integer a_i of A onto a line segment $y = a_i x + b_i$, such that all line segments passe through a point p, in Q. Solve the Minimum Cone Segment Cover problem. By the previous theorem, any solution will contain n segments, which will be given in the form of a sequence of triplets. In linear time extract from each triplet the successor of the point, who was mapped to the middle line equation in the triplet. □

3 The intersecting case

The algorithm computes all intersection points among the edges of E as follows : it sorts separately the set of vertices V_r lying on ray r, and the set of vertices $V_{r'}$, lying on ray r' , based on the distance of the vertices from O.

Let $e_{i,j} = (v_{r,i}, v_{r',j})$ be an edge of E, with vertex $v_{r,i}$ lying on r having rank i in the sorting of V_r; $v_{r',j}$ is defined in a similar way.

The algorithm intersects $e_{i,j}$ with $e_{i,k}$, $i = 1, \ldots, n$, $k = 1, \ldots, j-1$. Let T be the computed set of intersection points. The set T is then sorted radially around O. Let q_1 be the first intersection point in the sorted order. The algorithm constructs the subcone C_1 of C, determined by r and the ray r_1 passing through q_1. It then intersects C_1 with E, and obtain the set E_1 consisting of non-intersecting segments. The next step is a call to the algorithm for the non-intersecting case with C_1 and E_1 as parameters and which outputs a minimum cover for O given C_1 and E_1. The process is repeated over each consecutive sub-cone. After processing the last sub-cone, the minimum covers for the various sub-cones are merged to form the minimum cover for O, given C and E.

The correctness of this divide-and-conquer type of algorithm is based on the correctness of the algorithm that solves the non-intersecting case and on the following theorem.

Theorem 3.1 *Let M_1 and M_2 be two minimum covers for O, given E_i and C_i as defined above, $i = 1, 2$. Let $C_{1,2} = C_1 \cup C_2$ and $E_{1,2} = E_1 \cup E_2$. Denote by $M_{1,2}$ the minimum cover for O given $C_{1,2}$ and $E_{1,2}$, then $M_{1,2} = M_1 \cup M_2$.*

PROOF Immediate, by contradiction. □

The time analysis of this algorithm is straightforward : finding all intersection points requires $O(m)$, where m is the number of intersection points. This is due to the fact that if $e_{i,j}$ intersects with $e_{k,l}$, then $(v_{r,i}, v_{r,k})$ has the reverse order of $(v_{r',j}, v_{r',l})$. Sorting the set Q takes $m \log m$, solving the problem over a given sub-cone takes $O(n \log n)$. Finally merging all the contiguous minimum covers requires $O(m.n)$.

References

[B83] M. Ben-Or, *Lower Bounds for Algebraic Computation Trees*, Proceeding 15th Annual ACM Symposium on the Theory of Computation, 1983, 80-86

[DLS88] J. Dean, A. Lingas, and J.-R. Sack, *Recognizing Polygons, or How to Spy*, The Visual Computer **3**, 1988, 344-355.

[H85] P. Hanrahan, *Topological Shape Models*, PhD Thesis, University of Wisconsin-Madison, 1985.

[HV49] A. Horn and F. Valentine, *Some Properties of L-Sets in the Plane*, Duke Mathematics J. **16**, 1949, 131-140.

[LL84] C. Lee and D. Lee, *On a Circle-Cover Minimization Problem*, Information Processing Letters **18**, 1984, 109-115.

[M81] Kurt Mehlhorn, *Data Structures and Algorithm* **1**: *Sorting and Searching*, EATCS Monographs on Theoretical Computer Science, Springer-Verlag, 1981, 68-71.

[O87] J. O'Rourke, *Art Gallery Theorems and Algorithms*, Oxford University Press, Oxford, 1987.

[S85] R. Seidel, *A Method for Proving Lower Bounds for Certain Geometric Problems*, Computational Geometry, Ed G.T. Toussaint, Elseiver Science Publishers B.V. (North-Holland), 1985, 319-333.

[V53] F. Valentine, *Minimal Sets of Visibility*, Proceedings of the American Mathematical Society **4**, 1953, 917-921.