

An Approximation Scheme for Terrain Guarding

Matt Gibson, Gaurav Kanade, Erik Krohn, and Kasturi Varadarajan*

Department of Computer Science
University of Iowa
Iowa City, IA 52242-1419, USA
{mrgibson, gkanade, ekrohn, kvaradar}@cs.uiowa.edu

Abstract. We obtain a polynomial time approximation scheme for the terrain guarding problem improving upon several recent constant factor approximations. Our algorithm is a local search algorithm inspired by the recent results of Chan and Har-Peled [2] and Mustafa and Ray [15]. Our key contribution is to show the existence of a planar graph that appropriately relates the local and global optimum.

1 Introduction

A *1.5D terrain* is a polygonal chain in the plane that is x -monotone, that is, any vertical line intersects the chain at most once. A terrain T consists of a set of m vertices $\{v_1, v_2, \dots, v_m\}$. The vertices are ordered in increasing order with respect to their x -coordinates. There is an edge connecting v_i with v_{i+1} for all $i = 1, 2, \dots, m - 1$. For any two points $a, b \in T$, we say that a sees b if the line segment \overline{ab} lies entirely above or on the terrain.

In this paper, we consider the *discrete terrain guarding problem* in which we are given a terrain T and finite sets $X, G \subseteq T$. For a set $G' \subseteq G$, we say that G' covers/sees/guards X if every point in X can be seen by at least one point in G' . The goal of the problem is to find a minimum cardinality subset of G that covers X .

The motivation for guarding terrains comes from placing street lights or security sensors along roads, as well as constructing line-of-sight networks for radio broadcasting and other communication networks [1].

A closely related problem is the *art gallery problem*. Again the goal is to find a minimum cardinality guarding set, but in this setting we must guard a simple polygon. The basic version of this problem, *vertex guarding*, requires guards to be placed at the vertices of the polygon. Another version, *point guarding*, allows guards to be placed anywhere inside the polygon.

The art gallery problem was shown to be NP-complete by Lee and Lin [14] and was later shown to be APX-hard by Eidenbenz [7]. This means that there is an $\epsilon > 0$ such that no polynomial time algorithm can compute a guarding set whose cardinality is within a $(1 + \epsilon)$ factor of the cardinality of an optimal guarding set, unless $P = NP$. Ghosh gives an $O(\log n)$ -approximation algorithm

* Partially supported by NSF CAREER award CCR 0237431.

for vertex guarding an n -vertex simple polygon [11]. The point guarding problem seems to be much more difficult as not as much is known about it [5]. A constant factor approximation is given by Nilsson for the special case of the problem when the polygon is x -monotone [16]. Based on his result, Nilsson gives an $O(OPT^2)$ -approximation algorithm for rectilinear polygons.

Previous Work on Terrain Guarding. Chen et al. [3] claimed that terrain guarding is NP-hard, but the proof was never completed formally [12]. Most of the past research has gone into developing approximation algorithms. The first constant factor approximation was a combinatorial algorithm given by Ben-Moshe et al. [1]. Clarkson and Varadarajan [4] also give a constant factor approximation based on rounding a linear programming relaxation. King gave a simple combinatorial 4-approximation which was later determined to actually be a 5-approximation [12]. Recently, Elbassioni et al. [8] gave a 4-approximation that also works for the weighted case.

All of the approximation algorithms use the following “order claim”:

Claim. Let a, b, c, d be four points on the terrain in increasing order according to x -coordinate. If a sees c and b sees d then a sees d .

Natural attempts at constructing NP-hardness reductions for the terrain guarding problem do not work because of the order claim. Recently Krohn and King [13] were able to get around the order claim and prove NP-hardness.

Our Contribution. We give a polynomial time algorithm that returns a guard cover whose cardinality is at most $(1+\epsilon)\cdot OPT$ for any $\epsilon > 0$. Here, OPT denotes the cardinality of an optimal guard cover. Thus we obtain the first PTAS for the problem improving upon several recent constant factor approximations. Given the hardness result [13], this settles the computational complexity of the problem.

The inspiration for our work comes from the recent results of Chan and Har-Peled [2] and Mustafa and Ray [15]. Chan and Har-Peled show that a local search algorithm actually yields a PTAS for the maximum independent set problem given a collection of disks. Unlike a previous PTAS for the problem [9], their analysis does not use packing arguments and thus also applies to “pseudo-disks”. Mustafa and Ray consider several geometric hitting set and set cover problems and describe local search algorithms that yield PTASs. For instance, in a rather surprising result they obtain a PTAS for the problem of covering a set of points by the smallest number of a given set of disks. Both papers use separator theorems for planar graphs. In particular, they show that there exists a planar graph that relates the locally optimal solution returned by the local search and the global optimal solution. The separator theorem is then used to show that the locally optimal solution is not too much worse than the global optimum.

Our PTAS for the terrain guarding problem is also based on local search. Our key contribution is to show the existence of an appropriate planar graph even for the terrain guarding context. Having shown this, the rest of the analysis is very similar to that of Mustafa and Ray [15].

2 Guarding Terrains via Local Search

Recall that our input is a polygonal terrain, a set X of points on the terrain that need to be guarded, a set G of possible guard locations, and a parameter $0 < \epsilon < 1$. For purposes of exposition, we will initially assume that $X \cap G = \emptyset$. We later show how this assumption can be removed. We describe a polynomial time algorithm that returns a subset $Q \subseteq G$ that sees X , so that $|Q|$ is at most a factor $(1 + \epsilon)$ times the size of the smallest subset of G that sees X . Let n denote the input size – the number of vertices in the terrain, plus $|G|$, plus $|X|$.

We say that a subset of G that sees X is b -locally optimal if one cannot obtain a smaller set of guards that sees X by deleting at most b guards from it and inserting at most $b - 1$ guards.

Our algorithm simply returns a b -locally optimal solution for $b = \frac{\alpha}{\epsilon^2}$, where α is a suitably large constant, by performing local search. We start with some arbitrary $Q \subseteq G$ that covers X . For every subset $S \subseteq Q$ of size at most b , we see if there exists a subset $T \subseteq G \setminus Q$ of size at most $|S| - 1$ such that $(Q \setminus S) \cup T$ guards X . If so, we set $Q \leftarrow (Q \setminus S) \cup T$. Every such exchange decreases the size of Q by at least one, and as such can happen at most n times. Since there are $\binom{n}{b}$ subsets S to consider, the running time is bounded by $n^{O(b)}$.

2.1 Approximation Analysis

Let R' denote the optimal cover for X , and B' the set of guards output by our local search algorithm on termination. We show that $|B' \setminus R'| \leq (1 + \epsilon)|R' \setminus B'|$, and thus $|B'| \leq (1 + \epsilon)|R'|$. Let $R \equiv R' \setminus B'$, $B \equiv B' \setminus R'$, and abusing notation, let X denote the set after removing all points seen by $R' \cap B'$. So now both R and B cover X and we wish to show that $|B| \leq (1 + \epsilon)|R|$. We will refer to points in B as blue points and points in R as red points.

The following lemma is our main contribution; it shows that the *locality condition* of Mustafa and Ray [15] is satisfied.

Lemma 1. *There exists a planar graph $\mathcal{G} = (V \equiv R \cup B, E)$ with the property that for each $x \in X$, there is an edge (r, b) in \mathcal{G} between guards $r \in R$ and $b \in B$ that both see x .*

Before giving the proof, we show how the lemma implies that $|B| \leq (1 + \epsilon)|R|$; this is similar to [2, 15]. We need the following partition theorem on planar graphs due to Frederickson [10]. For $U \subseteq V$, let $\Gamma(U)$ denote the set of neighbors in \mathcal{G} of vertices in U with U excluded. Let $\mu = |V|$.

Lemma 2. *For any parameter $1 \leq r \leq \mu$, we can find a set $S \subseteq V$ of size at most $c_1\mu/\sqrt{r}$ and a partition of $V \setminus S$ into μ/r sets $V_1, V_2, \dots, V_{\mu/r}$, satisfying (i) $|V_i| \leq c_2r$, (ii) $|\Gamma(V_i)| \leq c_3\sqrt{r}$, and (iii) $(V_i \cup \Gamma(V_i)) \cap V_j = \emptyset$ for $i \neq j$. Here, c_1, c_2 , and c_3 are absolute positive constants.*

Let us apply the lemma with $r \equiv b/(c_2 + c_3)$. We have $|V_i \cup \Gamma(V_i)| \leq c_2r + c_3\sqrt{r} \leq b$. Thus, letting $R_i = R \cap V_i$ and $B_i = B \cap V_i$, we must have $|B_i| \leq$

$|R_i| + |\Gamma(V_i)|$. For otherwise, the local search can replace B_i by $R_i \cup \Gamma(V_i)$ and obtain a smaller set that still covers X (Lemma 1), a contradiction.

Thus

$$\begin{aligned} |B| &\leq |S| + \sum_i |B_i| \leq |S| + \sum_i |R_i| + \sum_i |\Gamma(V_i)| \leq |R| + c \frac{\mu}{\sqrt{r}} \\ &\leq |R| + c' \frac{|R| + |B|}{\sqrt{b}}, \end{aligned}$$

where c and c' are positive constants. With b a large enough constant times $1/\epsilon^2$, this implies that $|B| \leq (1 + \epsilon)|R|$.

2.2 Proof of Lemma 1

We begin with some notation. For points a and b on the terrain, we say $a \leq b$ to mean that the x -coordinate $a.x$ of a is at most $b.x$. We use the notation of intervals that this implies – for instance, $[a, b]$ denotes all points c on the terrain so that $a \leq c \leq b$.

We now prove Lemma 1. Let us first construct the planar graph \mathcal{G} . For each $x \in X$, let $\lambda(x)$ denote the leftmost point that sees x among points in $R \cup B$ to the left of x , assuming such a point does exist. Similarly, let $\rho(x)$ denote the rightmost point that sees x among points in $R \cup B$ to the right of x , assuming such a point does exist. Note that at least one of $\lambda(x)$ or $\rho(x)$ does exist.

Let A_1 denote the set of segments $\overline{\lambda(x)x}$, for $x \in X$. Because of the order claim, these segments do not cross. For each $v \in R \cup B$, shoot a vertical ray up from v ; if this ray hits some segment in A_1 , let $\overline{\lambda(y)y}$ denote the first such segment hit; we add the edge $(v, \lambda(y))$ to a set E_1 if v and $\lambda(y)$ are of opposite colors.

Now, the edges in $A_1 \cup E_1$ can be embedded above the terrain in a non-crossing way. To see this, let A_1 be embedded as the original straight line segments. To embed an edge $(v, \lambda(y)) \in E_1$ as above, we travel straight up from v till we hit $\overline{\lambda(y)y}$, and then slide along the segment $\overline{\lambda(y)y}$ to reach $\lambda(y)$. See Figure 1. A more formal argument that $A_1 \cup E_1$ can be so embedded is given in the appendix.

Let A_2 denote the set of segments $\overline{x\rho(x)}$, for $x \in X$. Again, these segments do not cross. For each $v \in R \cup B$, shoot a vertical ray up from v ; if this ray hits some segment in A_2 , let $\overline{y\rho(y)}$ denote the first such segment hit; we add the edge $(v, \rho(y))$ to a set E_2 if v and $\rho(y)$ are of opposite colors.

The edges in $A_2 \cup E_2$ can also be embedded above the terrain in a non-crossing way. We “flip” the embedding of $A_1 \cup E_1$ to obtain a non-crossing embedding below the terrain; see Figure 2. This gives us a planar embedding of $A_1 \cup E_1 \cup A_2 \cup E_2$.

Finally, for each $x \in X$, we add the edge $(\lambda(x), \rho(x))$ to a set E_3 if $\lambda(x)$ and $\rho(x)$ are of opposite colors. Our graph \mathcal{G} consists of the edge set $E_1 \cup E_2 \cup E_3$. This is a planar graph; just embed E_1 and E_2 as above, and for each $(\lambda(x), \rho(x)) \in E_3$, embed it using the embedding of the segments $\overline{\lambda(x)x}$ and $\overline{x\rho(x)}$.

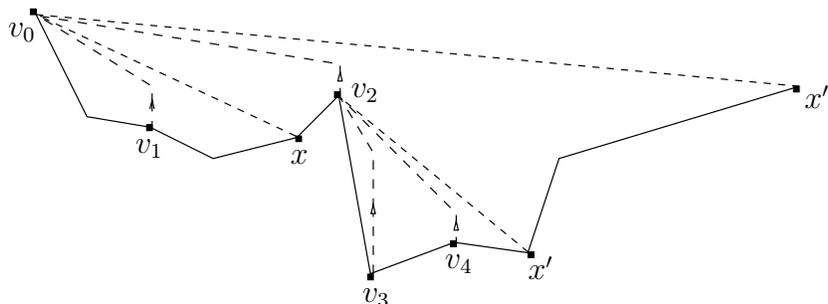


Fig. 1. The embedding of $A_1 \cup E_1$, with $X = \{x, x', x''\}$, and $R \cup B = \{v_0, v_1, v_2, v_3, v_4\}$. Segments in A_1 are shown in dashed lines, and the edges in E_1 are embedded as dashed curves with arrows. Note that $v_0 = \lambda(x) = \lambda(x'')$, and $v_2 = \lambda(x')$.

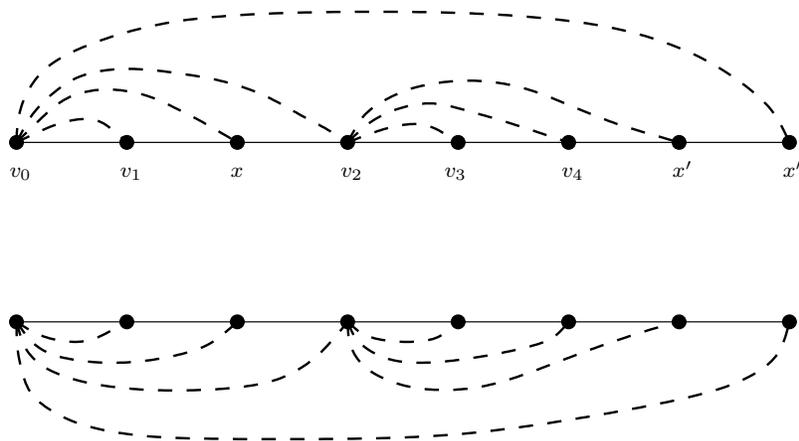


Fig. 2. A combinatorial embedding of $A_1 \cup E_1$ from Figure 1, and flipping it so that $A_1 \cup E_1$ is now embedded below the terrain. Note that only the edges in $A_1 \cup E_1$ are being flipped to make room for $A_2 \cup E_2$; the vertex set $R \cup B \cup X$ retains its embedding.

Now we need to show that for each $x \in X$, there are points $r \in R$ and $b \in B$ that see x , and $(r, b) \in E_1 \cup E_2 \cup E_3$. Fix an $x \in X$. If $\lambda(x)$ and $\rho(x)$ are of opposite colors, then $(\lambda(x), \rho(x)) \in E_3$, and we are done. Otherwise, it must be the case that there are red and blue points to the left of x that see x , or that there are red and blue points to the right of x that see x .

Let us assume that the first case holds (there are red and blue points to the left of x that see x), and that $\lambda(x)$ is red. The other situations are symmetric. Let b be the leftmost blue point that sees x ; it must be that $b \in (\lambda(x), x)$. Thus the ray shot up from b hits $\overline{\lambda(x)x}$; let $\lambda(y)y$ be the first segment in A_1 that it hits. See Figure 3. Because segments in A_1 don't cross, it must be that $\lambda(y) \in [\lambda(x), b)$ and $y \in (b, x)$. The order claim (applied to $\lambda(y)$, b , y , and x) implies that $\lambda(y)$ sees x . Now $\lambda(y)$ cannot be blue, otherwise b is not the leftmost blue point that sees x . Thus $\lambda(y) \in R$, $b \in B$, both $\lambda(y)$ and b see x and $(b, \lambda(y)) \in E_1$. This completes the proof.

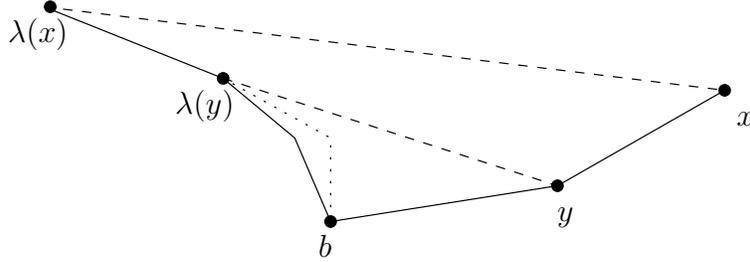


Fig. 3. Here, $\lambda(y)$ sees y and b sees x . By the order claim, $\lambda(y)$ sees x .

2.3 Relaxing the Disjointness Assumption

For ease of exposition, we have so far assumed that the set of possible guard locations has an empty intersection with the set of points to be guarded. We now relax that assumption. The algorithm remains unchanged, and we indicate how the analysis is modified. For each $x \in X$, the point $\lambda(x)$ (resp. $\rho(x)$) denotes the leftmost (resp. rightmost) point that sees x among points in $R \cup B$ *strictly* to the left (resp. right) of x , assuming such a point does exist. With this understanding, the construction of the planar graph \mathcal{G} proceeds with the sets A_1 , E_1 , A_2 , and E_2 defined exactly as above.

There is a change in the construction of E_3 . For each $x \in X$ that is not in $R \cup B$, we proceed as before and add the edge $(\lambda(x), \rho(x))$ to E_3 if $\lambda(x)$ and $\rho(x)$ are of opposite colors. For $x \in X$ that is also in $R \cup B$, we add the edge $(\lambda(x), x)$ to E_3 if $\lambda(x)$ exists and $\lambda(x)$ and x are of opposite colors; we also add the edge $(x, \rho(x))$ to E_3 if $\rho(x)$ exists and $\rho(x)$ and x are of opposite colors.

This completes the construction of the graph, which is readily seen to be planar. To show Lemma 1, we need to argue that for each $x \in X$, there are points $r \in R$ and $b \in B$ that see x , and $(r, b) \in E_1 \cup E_2 \cup E_3$. If $x \notin R \cup B$ then the argument is exactly as before. Without loss of generality, assume that $x \in R$. Since B sees X , there must be a point in B that sees x . Assume that such a blue point lies to the left of x . The other case is symmetric. In this case, it must be that $\lambda(x)$ exists. If $\lambda(x) \in B$, then we are done since we added $(\lambda(x), x)$ to E_3 . Therefore assume that $\lambda(x) \in R$. Let b be the leftmost blue point that sees x . Now we can use the reasoning in the last paragraph of the previous section to show that there is an edge $(b, u) \in E_1$ such that both b and u see x , b is blue, and u is red.

3 Conclusions

We have shown that the discrete terrain guarding problem admits a polynomial-time approximation scheme. We can also obtain a PTAS in the scenario where the possible guard locations are from a finite set, and we want to see the entire terrain – this problem is readily reduced to discrete terrain guarding. In the continuous terrain guarding problem, guards are allowed to be located anywhere on the terrain. The local search can be seen to work even here, and we can show that a single iteration of the local search can be implemented in polynomial time along the lines of Section 4 of [6]. There is one issue that remains, however, and this is to bound the number of bits needed to represent the guards maintained by the local search. We are currently investigating how this can be handled.

Acknowledgements. We thank the anonymous reviewers for useful feedback.

References

1. Boaz Ben-Moshe, Matthew J. Katz, and Joseph S. B. Mitchell. A constant-factor approximation algorithm for optimal terrain guarding. In *SODA*, pages 515–524, 2005.
2. Timothy Chan and Sarel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. In *Symposium on Computational Geometry*, 2009. To Appear.
3. Danny Z. Chen, Vladimir Estivill-Castro, and Jorge Urrutia. Optimal guarding of polygons and monotone chains (extended abstract), 1996.
4. Kenneth L. Clarkson and Kasturi Varadarajan. Improved approximation algorithms for geometric set cover. In *SCG '05: Proceedings of the twenty-first annual symposium on Computational geometry*, pages 135–141, New York, NY, USA, 2005. ACM.
5. Ajay Deshpande, Taejung Kim, Erik D. Demaine, and Sanjay E. Sarma. A pseudopolynomial time $o(\log^2 n)$ -approximation algorithm for art gallery problems. In *Proceedings of the 10th Workshop on Algorithms and Data Structures (WADS 2007)*, volume 4619 of *Lecture Notes in Computer Science*, pages 163–174, Halifax, Nova Scotia, Canada, August 15–17 2007.

6. Alon Efrat and Sarel Har-Peled. Guarding galleries and terrains. *Information Processing Letters*, 100(6):238 – 245, 2006.
7. Stephan Eidenbenz. Inapproximability results for guarding polygons without holes. In *Lecture Notes in Computer Science*, pages 427–436. Springer, 1998.
8. Khaled Elbassioni, Erik Krohn, Domagoj Matijevec, Julian Mestre, and Domagoj Severdija. Improved approximations for guarding 1.5-dimensional terrains. In Susanne Albers and Jean-Yves Marion, editors, *26th International Symposium on Theoretical Aspects of Computer Science (STACS 2009)*, Dagstuhl, Germany, 2009. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany.
9. Thomas Erlebach, Klaus Jansen, and Eike Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM Journal on Computing*, 34(6):1302–1323, 2005.
10. Greg N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Comput.*, 16(6):1004–1022, 1987.
11. S. Ghosh. Approximation algorithms for art gallery problems, *Proc. Canadian Information Processing Society Congress*, 1987.
12. James King. A 4-approximation algorithm for guarding 1.5-dimensional terrains. In *Proc. of the 7th Latin American Symposium on Theoretical Informatics (LATIN'06)*, pages 629–640, Valdivia, Chile, 2006.
13. Erik Krohn and James King. The complexity of guarding terrains, *Manuscript*, 2009.
14. D. Lee and A. Lin. Computational complexity of art gallery problems. *Information Theory, IEEE Transactions on*, 32(2):276–282, Mar 1986.
15. Nabil H. Mustafa and Saurabh Ray. Ptas for geometric hitting set problems via local search. In *Symposium on Computational Geometry*, 2009. To Appear.
16. Bengt J. Nilsson. Approximate guarding of monotone and rectilinear polygons. In *ICALP*, pages 1362–1373, 2005.

A Appendix

We present here a more formal argument for a piece of Lemma 1 that shows that the graph $(X \cup R \cup B, A_1 \cup E_1)$ has a planar embedding with $X \cup R \cup B$ on a horizontal line and $A_1 \cup E_1$ drawn above the line. We now think of A_1 as a set of combinatorial edges rather than as line segments. Let us say that two edges in $A_1 \cup E_1$ *cross* if all four end points are distinct, and can be ordered as $a < b < c < d$ with the edges being (a, c) and (b, d) . Notice that crossing in this sense is determined entirely by the ordering of the endpoints, and makes no reference to a drawing of the edges. We first argue that no two edges in $A_1 \cup E_1$ cross.

1. Let $(\lambda(x), x)$ and $(\lambda(y), y)$ be any two edges in A_1 . They do not cross, for if $\lambda(x) < \lambda(y) < x < y$, then by the order claim, y sees $\lambda(x)$, contradicting the definition of $\lambda(y)$.
2. We argue that an edge in A_1 and an edge in E_1 do not cross. Let us recall how edges in E_1 are defined. Let us say that an edge $(\lambda(x), x)$ is *above* $v \in R \cup B$ if $\lambda(x) < v < x$. We look at all edges in A_1 that are above v , and if this set is non-empty we find the “innermost” such edge $(\lambda(y), y)$. We add the

edge $(\lambda(y), v)$ to E_1 if $\lambda(y)$ and v are of opposite colors. Let us say that $(\lambda(y), y) \in A_1$ defines $(\lambda(y), v)$ in this case.

Now suppose $(\lambda(y), v)$ crosses some $(\lambda(x), x) \in A_1$. Now if $\lambda(x) < \lambda(y) < x < v$, then $(\lambda(x), x)$ and $(\lambda(y), y)$ cross, a contradiction. On the other hand, if $\lambda(y) < \lambda(x) < v < x$, there are two cases: either $x \leq y$, in which case the contradiction is that $(\lambda(y), y)$ is not the innermost edge in A_1 that is above v ; or $x > y$, in which case the contradiction is that $(\lambda(y), y)$ and $(\lambda(x), x)$, which are both edges in A_1 , cross.

3. We now show that no two edges in E_1 cross. Consider two such edges $(\lambda(y_1), v_1)$ and $(\lambda(y_2), v_2)$, defined by $(\lambda(y_1), y_1)$ and $(\lambda(y_2), y_2)$ respectively. These edges do not cross, for if $\lambda(y_1) < \lambda(y_2) < v_1 < v_2$, then $(\lambda(y_1), v_1)$ crosses $(\lambda(y_2), y_2)$, which contradicts the fact that an edge in E_1 and an edge in A_1 do not cross.

Thus, no two edges in $A_1 \cup E_1$ cross. From this, it follows that the required embedding exists. For instance, embed the vertices $X \cup B \cup R$ as distinct points in the correct order on the lower half of the unit circle, and draw the edges in $A_1 \cup E_1$ as straight line segments. Using convexity arguments, this can be seen to be a planar embedding. Now we “bend” the lower half of the unit circle into a horizontal segment, allowing the drawing of edges in $A_1 \cup E_1$ to now become curved.