Geometric Hitting Set Problems

K Prahlad Narasimhan

National Institute of Science Education and Research,
HBNI, Bhubaneswar, India

November, 2021

1 Introduction

In this report we present a local search algorithm for geometric hitting set problems that was obtained by Mustafa and Ray in [7]. We then look at three applications of this algorithm - hitting halfspaces in $\mathbb{R}^3$, dominating terrain-like graphs and guarding weakly visible polygons. Finally, we prove the optimality of local search algorithms for certain geometric hitting set problems. In this introductory section, we first define the Hitting Set problem and prove that it is equivalent to the Set Cover problem. This is done in Lemma 1.1. We then present a simple family of local search algorithms for the Hitting Set problem (Algorithm 1.2) and compute its running time. The main result of this paper is to prove that this family of algorithms is a PTAS for a restricted case of the Hitting Set problem. We take a small detour and prove *Radon's Theorem* as Theorem 1.5 before proving this result in Section 2.

For a set P, we let $\mathcal{P}(P)$ denote the power set of P.

Problem (Hitting Set). Let P be a finite set and $\mathcal{D} \subseteq \mathcal{P}(P)$. A set $I \subseteq P$ is called a *hitting set* of $\mathcal{D}$ if $I \cap D \neq \emptyset$ for all $D \in \mathcal{D}$. Find a hitting set I' of $\mathcal{D}$ of the smallest size.

The tuple $(P, \mathcal{D})$ in an instance of the Hitting Set problem is denoted by $\mathcal{R}(P, \mathcal{D})$ and is called the *range space* of this instance. P is called the *ground set* and $\mathcal{D}$ is called the *set of ranges*.

Problem (Set Cover). Let E be a finite set of elements and $\mathcal{S} \subseteq \mathcal{P}(E)$. A set $J \subseteq \mathcal{S}$ is called a *set cover* of E if $\cup_{S \in J} S = E$. Find a set cover J' of E of the smallest size.

Lemma 1.1. The Hitting Set and Set Cover problems are equivalent.

Proof. Let $\mathcal{R}(P, \mathcal{D})$ be an instance of the Hitting Set problem. Let $P = \{p_1, p_2 \dots p_n\}$ and $\mathcal{D} = \{D_1, D_2 \dots D_m\}$. Construct a Set Cover instance $(E, \mathcal{S})$ as follows: let $E = \mathcal{D}$ and $\mathcal{S} = \{S_1, S_2 \dots S_n\}$ where S_i is defined to be $\{D_j \in \mathcal{D} \mid p_i \in D_j\}$ for each i. Let $I = \{p_{i_1}, p_{i_2} \dots p_{i_k}\} \subseteq P$ be a hitting set of $\mathcal{D}$. Consider the corresponding subset $J = \{S_{i_1}, S_{i_2} \dots S_{i_k}\}$ of $\mathcal{S}$. We prove that J is a set cover of E. Let $e \in E$ be an arbitrary element. Since $E = \mathcal{D}$, $e = D_j$ for some j. Since I is a hitting set of $\mathcal{D}$, there is a $p_{i_l} \in I$ such that $p_{i_l} \in D_j$. By construction, $e = D_j \in S_{i_l}$. Thus, J covers e. Since e was arbitrary, J is a set cover of E. Using symmetric arguments, it is easy to see that a set cover of E in the constructed instance corresponds to a hitting set of $\mathcal{D}$ of the original instance. This implies that the hitting sets of $\mathcal{R}(P, \mathcal{D})$ are in one-one correspondence with the set covers of $(E, \mathcal{S})$. Since this correspondence also preserves cardinality, an optimal solution of the Hitting Set instance corresponds to one for the Set Cover instance.

Now, consider a Set Cover instance $(E, \mathcal{S})$ where $E = \{e_1, e_2 \dots e_n\}$ and $\mathcal{S} = \{S_1, S_2 \dots S_m\}$. Consider the Hitting Set instance $\mathcal{R}(P, \mathcal{D})$ where $P = \mathcal{S}$ and $\mathcal{D} = \{D_1, D_2 \dots D_n\}$ where D_i is defined to be $\{S_j \mid e_i \in S_j\}$ for each i. Clearly, a set cover of $(E, \mathcal{S})$ is a hitting set of $\mathcal{R}(P, \mathcal{D})$ and the vice versa. This completes the proof of this lemma since any algorithm that approximates the Set Cover problem (respectively the Hitting Set problem) can be used to approximate, with the same factor, the Hitting Set problem (respectively the Set Cover problem). □

Algorithm 1.2 (k-Level Local Search). Let k be a fixed natural number and $\mathcal{R}(P, \mathcal{D})$ be an instance of the Hitting Set problem. We construct a hitting set S of $\mathcal{D}$ as follows: Let S be P. If there exists $U \subseteq S$ where $|U| = k$ and $U' \subseteq P$ where $|U'| = k - 1$ such that $(S \setminus U) \cup U'$ remains a hitting set, replace S by $(S \setminus U) \cup U'$. Repeat this procedure till no such U and U' exist.

Lemma 1.3. Given an input instance $\mathcal{R}(P, \mathcal{D})$ of the Hitting Set problem with $|P| = n$ and $|\mathcal{D}| = m$, Algorithm 1.2 runs in $\mathcal{O}(mn^{2k+1})$ time.

Proof. Clearly, at each improvement step the size of S decreases by one. Thus, Algorithm 1.2 runs at most n times. At each step, we need to check at most ${}^nC_k \cdot {}^nC_{k-1}$ sets to choose U and U' or determine that no such sets exist. Checking whether an improvement is a hitting set takes $\mathcal{O}(mn)$ time. Thus, the algorithm runs in $\mathcal{O}(mn^{2k+1})$ time. □

In the next section (Lemma 2.5) we will prove the following: given a range space $\mathcal{R}(P, \mathcal{D})$ that follows the *locality condition* (defined below) and an arbitrary $\epsilon > 0$, one can find a $k \in \mathcal{O}(\epsilon^{-2})$ in polynomial time such that the k-level local search algorithm returns a $(1 + \epsilon)$-approximate solution for any instance of the Hitting Set problem.

Definition (Locality Condition). A range space $\mathcal{R}(P, \mathcal{D})$ is said to satisfy the *locality condition* if, for any two disjoint subsets R and B of P, it is possible to construct a planar bipartite graph $G((R, B), E)$ - that is, a graph with bipartition R and B and edge set E, such that for any $D \in \mathcal{D}$ which intersects both R and B, there exists a $(u, v) \in E$ where $u \in D \cap R$ and $v \in D \cap B$.

In Section 3 we will prove that range spaces arising from some typical geometric hitting set problems satisfy the locality condition. We conclude this section by proving *Radon's Theorem.* First, we prove a small lemma which is referenced from these lecture notes [6].

Lemma 1.4. Let $n \in \mathbb{N}$ and $P = \{p_1, p_2 \dots p_{n+2}\} \subset \mathbb{R}^n$. Then, there exists $\{\alpha_1, \alpha_2 \dots \alpha_{n+2}\} \subset \mathbb{R}$ such that:

(i) There exists a $\alpha_i \in \{\alpha_1, \alpha_2 \dots \alpha_{n+2}\}$ which is non-zero.

(ii) $\sum_{i=1}^{n+2} \alpha_i p_i = 0$.

(iii) $\sum_{i=1}^{n+2} \alpha_i = 0$.

Proof. Consider the points $P' = \{p_2 - p_1, p_3 - p_1 \dots p_{n+2} - p_1\}$. Since $|P'| = n+1$, P' is linearly dependent.

Hence, there exists $\{\beta_2, \beta_3 \dots \beta_{n+2}\} \subset \mathbb{R}$, not all elements zero, such that $\sum_{i=2}^{n+2} \beta_i(p_i - p_1) = 0$. Hence,

$$\sum_{i=2}^{n+2} \beta_i p_i - \sum_{i=2}^{n+2} \beta_i \cdot p_1 = 0$$

Let $\alpha_1 = -\sum_{i=2}^{n+2} \beta_i$ and $\alpha_i = \beta_i$ for all $i \in \{2, 3 \dots n+2\}$. Then, clearly, $\{\alpha_1, \alpha_2 \dots \alpha_{n+2}\}$ satisfy the three required conditions. $\square$

Definition (Convex Hull). The *convex hull* of P, where $P \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}$, is the smallest convex set that contains P. This set is denoted by $\textsc{conv}(P)$.

Note that the convex hull of a set is indeed a convex set.

Theorem 1.5 (Radon's Theorem). Let $n \in \mathbb{N}$ and $P = \{p_1, p_2 \dots p_{n+2}\} \subset \mathbb{R}^n$. Then, there exists a partition P_1, P_2 of P such that $\textsc{conv}(P_1) \cap \textsc{conv}(P_2) \neq \emptyset$.

Proof. By Lemma 1.4, there exists a set $\{\alpha_1, \alpha_2 \dots \alpha_{n+2}\} \subset \mathbb{R}$, not all elements zero, such that $\sum_{i=1}^{n+2} \alpha_i p_i = 0$ and $\sum_{i=1}^{n+2} \alpha_i = 0$. Let $P_1 = \{p_i \mid \alpha_i > 0\}$ and $P_2 = \{p_i \mid \alpha_i \leq 0\}$. Let $S = \sum_{i:p_i \in P_1} \alpha_i$. Note that $S \neq 0$. Since $\textsc{conv}(P_1)$ is a convex set, $x = \sum_{i:p_i \in P_1} \frac{\alpha_i}{S} p_i \in \textsc{conv}(P_1)$. Furthermore, since $\sum_{i=1}^{n+2} \alpha_i = 0$, $S = -\sum_{i:p_i \in P_2} \alpha_i$. Hence,

$$x = \sum_{i:p_i \in P_1} \frac{\alpha_i}{S} p_i = \frac{1}{S} \cdot \sum_{i:p_i \in P_1} \alpha_i p_i = \frac{-1}{S} \cdot \sum_{i:p_i \in P_2} \alpha_i p_i = \sum_{i:p_i \in P_2} \frac{\alpha_i}{-S} p_i \in \textsc{conv}(P_2)$$

Hence, $x \in \textsc{conv}(P_1) \cap \textsc{conv}(P_2)$. This completes the proof of this theorem. $\square$

2 Approximations Using Local Search

In this section, we design a PTAS for the Hitting Set problem in range spaces that satisfy the locality condition. That is, we will prove that there exists a collection of algorithms that run in polynomial time, say $\{A_\epsilon\}_{\epsilon > 0}$, such that A_ϵ produces a $(1+\epsilon)$-approximate solution for Hitting Set. For any vertex v of a graph $G(V, E)$, we let $N_G(v) = \{u \in V \mid (u, v) \in E\}$. For a subset W of V, we define $N_G(W)$ as $\cup_{w \in W} N_G(w)$.

Lemma 2.1. Let $\mathcal{R}(P, \mathcal{D})$ be a range space that satisfies the locality condition. Let R be an optimal hitting set of $\mathcal{D}$ and B be the hitting set that is returned by a k-level local search algorithm. Assume that R and B are disjoint. Then there exists a planar bipartite graph $G((R, B), E)$ such that for any $B' \subseteq B$ with $|B'| \leq k$, $(B \setminus B') \cup N_G(B')$ is a hitting set of $\mathcal{D}$.

Proof. Let $G((R, B), E)$ be the planar bipartite graph that is ensured by $\mathcal{R}(P, \mathcal{D})$ satisfying the locality condition. Let $D \in \mathcal{D}$ be an arbitrary element. Since both R and B are hitting sets of $\mathcal{D}$, $D \cap R$ and $D \cap B$ are non-empty. Assume that D is only hit by elements of B'. Then, by the locality condition, there is a $u \in R \cap D$ and $v \in B \cap D = B' \cap D$ such that $(u, v) \in E$. Thus, $(B \setminus B') \cup N_G(B')$ hits D. If D is hit by elements of B outside B', then, clearly, $(B \setminus B') \cup N_G(B')$ hits D. Hence, $(B \setminus B') \cup N_G(B')$ is a hitting set of $\mathcal{D}$. $\square$

Corollary 2.2. Let $\mathcal{R}(P, \mathcal{D})$ be a range space that satisfies the locality condition. Let R be an optimal hitting set of $\mathcal{D}$ and B be the hitting set that is returned by a k-level local search algorithm. Assume that R and B are disjoint. Then there exists a planar bipartite graph $G((R, B), E)$ such that for any $B' \subseteq B$ with $|B'| \leq k$, $|N_G(B')| \geq |B'|$.

Proof. By Lemma 2.1, $(B \setminus B') \cup N_G(B')$ is a hitting set of $\mathcal{D}$. Since no local improvements are possible in B, $|N_G(B')| \geq |B'|$. □

Let $G(V, E)$ be a graph. Let $\mathcal{G} = \{G_1, G_2 \ldots G_m\}$ be a collection of subsets of V such that $\cup_{i=1}^m G_i = V$. For a $G_i \in \mathcal{G}$, define the *boundary* of G_i, denoted by $\mathcal{B}(G_i)$, as $\{v \mid v \in G_i \cap G_j \text{ for some } G_j \in \mathcal{G} \setminus \{G_i\}\}$. Define $\mathcal{I}(G_i)$, the *interior* of G_i, as $G_i \setminus \mathcal{B}(G_i)$ for each $G_i \in \mathcal{G}$. We use the following result proved by Frederickson in 1987 [3].

Theorem 2.3 (Planar Graph Partition). Let $H(V, E)$ be a planar graph and let $|V| = n$. Let t be an arbitrary natural number. Then, there exists a collection of subsets of vertices, $\mathcal{G} = \{G_1, G_2 \ldots G_m\}$ with $\cup_{i=1}^m G_i = V$ and $|G_i| \leq t$ for all $G_i \in \mathcal{G}$ such that $\sum_{i=1}^m |\mathcal{B}(G_i)| \leq \frac{\gamma n}{\sqrt{t}}$ where γ is some fixed constant.

Theorem 2.4. Let $G((R, B), E)$ be a planar bipartite graph where $|R| \geq 2$ such that for all $B' \subseteq B$ with $|B'| \leq k$, for some $k \in \mathbb{N}$, $|N_G(B')| \geq |B'|$. Then, $|B| \leq (1 + \frac{c}{\sqrt{k}})|R|$ for some $c \in \mathbb{R}$.

Proof. By Theorem 2.3, there exists a collection of subsets of vertices, $\mathcal{G} = \{G_1, G_2 \ldots G_m\}$ such that $\cup_{i=1}^m G_i = R \cup B$ and $|G_i| \leq k$ for all $G_i \in \mathcal{G}$. Furthermore, $\sum_{i=1}^m |\mathcal{B}(G_i)| \leq \frac{\gamma n}{\sqrt{k}}$ where $n = |R| + |B|$ and γ is some fixed constant. Let $\mathcal{B}^R(G_i)$ and $\mathcal{B}^B(G_i)$ be the boundary vertices of G_i in R and B respectively. $\mathcal{I}^R(G_i)$ and $\mathcal{I}^B(G_i)$ are defined similarly. Let $G_i \in \mathcal{G}$ be an arbitrary element. Since $|G_i| \leq k$, $|\mathcal{I}^B(G_i)| \leq k$. Hence, by our supposition,

$$|\mathcal{I}^B(G_i)| \leq |N_G(\mathcal{I}^B(G_i))|$$

Note that $N_G(\mathcal{I}^B(G_i)) \subseteq G_i \cap R$. This implies that $|\mathcal{I}^B(G_i)| \leq |\mathcal{I}^R(G_i)| + |\mathcal{B}^R(G_i)|$. Adding $|\mathcal{B}^B(G_i)|$ on both sides, and noting that $G_i \cap B = \mathcal{I}^B(G_i) \cup \mathcal{B}^B(G_i)$, we get:

$$|\mathcal{I}^B(G_i)| + |\mathcal{B}^B(G_i)| \leq |\mathcal{I}^R(G_i)| + |\mathcal{B}^R(G_i)| + |\mathcal{B}^B(G_i)| \implies |G_i \cap B| \leq |\mathcal{I}^R(G_i)| + |\mathcal{B}(G_i)|$$

Summing over all $G_i \in \mathcal{G}$,

$$\begin{aligned}
&\sum_{i=1}^m |G_i \cap B| \leq \sum_{i=1}^m |\mathcal{I}^R(G_i)| + \sum_{i=1}^m |\mathcal{B}(G_i)| \\
&\implies |B| \leq \sum_{i=1}^m |\mathcal{I}^R(G_i)| + \sum_{i=1}^m |\mathcal{B}(G_i)| && (\cup_{i=1}^m G_i = R \cup B) \\
&\implies |B| \leq \sum_{i=1}^m |\mathcal{I}^R(G_i)| + \frac{\gamma n}{\sqrt{k}} \\
&\implies |B| \leq |R| + \gamma \frac{|R| + |B|}{\sqrt{k}} \\
&\implies (1 - \frac{\gamma}{\sqrt{k}})|B| \leq (1 + \frac{\gamma}{\sqrt{k}})|R|
\end{aligned}$$

Fix $k \geq 4\gamma^2$ and set $c = 4\gamma$. Also, let $\alpha = \frac{\gamma}{\sqrt{k}}$. Note that $\alpha \leq \frac{1}{2}$. If we let $\mathbb{N}_0$ denote $\mathbb{N} \cup \{0\}$, we have:

$$\begin{aligned}
|B| &\leq \frac{1+\alpha}{1-\alpha}|R| \\
&= (1+\alpha) \cdot \sum_{i \in \mathbb{N}_0} \alpha^i \cdot |R| \\
&= (1+\alpha) \cdot (1 + \alpha + \sum_{i=2}^{\infty} \alpha^i) \cdot |R| \\
&= (1+\alpha) \cdot (1 + \alpha + \alpha \cdot \frac{\alpha}{1-\alpha}) \cdot |R| \\
&\leq (1+\alpha) \cdot (1 + \alpha + \alpha) \cdot |R| && (\alpha \leq 1 - \alpha) \\
&= (1 + 3\alpha + 2\alpha^2) \cdot |R| \\
&\leq (1 + 4\alpha) \cdot |R| && (2\alpha^2 \leq \alpha) \\
&= (1 + \frac{c}{\sqrt{k}}) \cdot |R|
\end{aligned}$$

Hence, $|B| \leq (1 + \frac{c}{\sqrt{k}})|R|$, proving our claim. □

Lemma 2.5. Let $\mathcal{R}(P, \mathcal{D})$ be a range space that satisfies the locality condition. Then, given an $\epsilon > 0$, there exists a $k \in \mathcal{O}(\epsilon^{-2})$ such that the k-level local search algorithm returns an $(1+\epsilon)$-approximate solution for this instance of the HITTING SET problem.

Proof. Let R be an optimum hitting set of $\mathcal{R}(P, \mathcal{D})$. Let γ be the constant specified in Theorem 2.3 and $c = 4\gamma$. Fix $k = \frac{c^2}{\epsilon^2}$. Let B be the hitting set returned by the k-level local search algorithm.

Case 2.5.1 ($R \cap B = \emptyset$). By Corollary 2.2, there is a planar bipartite graph $G((R, B), E)$ such that for any $B' \subseteq B$ where $|B'| \leq k$, $|N_G(B')| \geq |B'|$. Hence, by Theorem 2.4, $|B| \leq (1 + \frac{c}{\sqrt{k}})|R|$. Replacing k by $\frac{c^2}{\epsilon^2}$, we have our claim.

Case 2.5.2 ($R \cap B \neq \emptyset$). Let $\hat{R} = R \setminus (R \cap B)$ and $\hat{B} = B \setminus (R \cap B)$. Furthermore, we let $\hat{P} = P \setminus (R \cap B)$ and $\hat{\mathcal{D}} = \mathcal{D} \setminus \{D \in \mathcal{D} \mid R \cap B \text{ hits } D\}$. Clearly, $\mathcal{R}(\hat{P}, \hat{\mathcal{D}})$ satisfies the locality condition. Also, note that $\hat{R}$ is an optimal hitting set of $\mathcal{R}(\hat{P}, \hat{\mathcal{D}})$ and $\hat{B}$ is a k-level locally optimum hitting set of this instance. Hence, by Case 2.5.1, we obtain: $|\hat{B}| \leq (1+\epsilon) \cdot |\hat{R}|$. Adding $|R \cap B|$ both sides, we have,

$$|\hat{B}| + |R \cap B| \leq (1+\epsilon) \cdot |\hat{R}| + |R \cap B| \implies |B| \leq (1+\epsilon) \cdot (|\hat{R}| + |R \cap B|) \implies |B| \leq (1+\epsilon) \cdot |R|$$

This completes the proof of this case.

Hence, an $\mathcal{O}(\epsilon^{-2})$-level local search algorithm returns a $(1+\epsilon)$-approximate solution for this instance of the HITTING SET problem. □

In Section 4, we will prove that this algorithm is, up to sublinear factors, optimal.

3 Problems Satisfying the Locality Condition

In this section, we will design a PTAS for three geometric problems: hitting halfspaces in $\mathbb{R}^3$, dominating *terrain-like graphs* and guarding *weakly visible* polygons. We do this by proving that each of these problems satisfy the locality condition. Then, given any $\epsilon > 0$, we run Algorithm 1.2 on the appropriate k given by Lemma 2.5 to get a $(1 + \epsilon)$-approximate solution.

3.1 Hitting Halfspaces

Let $X \subseteq \mathbb{R}^3$ be a finite set of points in and $\mathcal{H}$ be a finite set of open halfspaces of $\mathbb{R}^3$. We are required to find a subset of X of minimum size which has at least one point from all the halfspaces of $\mathcal{H}$. We guess an $o \in X$ and add it to our solution. Let $P = X \setminus \{o\}$ and $\mathcal{D} \subseteq \mathcal{H}$ be the subset of halfspaces that are not incident on o. Let $\mathcal{R}(P, \mathcal{D})$ be the range space corresponding to P and $\mathcal{D}$. Note the abuse of notation here: since the set of ranges is a subset of the power set of the ground set, we should have first taken the intersection of the halfspaces of $\mathcal{H}$ with P before defining $\mathcal{R}(P, \mathcal{D})$. However, since this notation is more convenient, we use it in this subsection.

If we produce S, a near optimal hitting set of $\mathcal{R}(P, \mathcal{D})$, then $S \cup \{o\}$ is clearly a near optimal hitting set of $\mathcal{R}(X, \mathcal{H})$. We now prove that $\mathcal{R}(P, \mathcal{D})$ satisfies the locality condition.

Theorem 3.1. Let R and B be two disjoint subsets of P. Then, there exists a planar bipartite graph $G((R, B), E)$ with the following property: given any open halfspace H of $\mathcal{H}$ such that $H \cap R$ and $H \cap B$ are non-empty, there is a $u \in H \cap R$ and a $v \in H \cap B$ such that $(u, v) \in E$.

Proof. We construct this graph by embedding the the points of $R \cup B$ onto the faces of the convex hull, say C, of this set. We ensure that the edges of E are within each face and are non-crossing. Once we have this, we choose a point s which is sufficiently close to one of the faces of C and a plane parallel to this face and lies below C. We map each point p embedded in the convex hull to the point that $\overrightarrow{sp}$ hits the plane. It is clear that this embedding onto the plane has non-crossing edges and hence proves that G is planar.

We refer to the points in R as *red* points and those in B as *blue* points. We add edges to E in two steps. In step one, we add all the red-blue edges (edges which are incident on one red point and one blue point) of C to E. In preparation for the second step, we triangulate the faces of C. For each $p \in R \cup B$ which is not a vertex of C, we let $\pi(p)$ be the point on the face of C that $\overrightarrow{op}$ intersects. For the sake of simplicity, we assume that none of the points map to the edges of the triangulated faces of C. Since $\mathcal{H}$ is a set of open halfspaces, this assumption is indeed valid since we can slightly "perturb" the points of an instance so that the corresponding range space instance does not change. Let Δ denote a triangulated face of C and Q denote the set of points mapped to to a point inside Δ. We have four cases based on the number of red vertices that Δ contains. We only need to consider the cases where Δ has 2 or 3 red vertices since the other two cases will follow using symmetrical arguments for blue vertices.

Case 3.1.1 (Δ has 2 red vertices and 1 blue vertex). We add edges between the projections of red points of Q and the blue vertex of Δ to E. We also add the edges between the projections of blue points of Q and the 2 red vertices of Δ. This is illustrated in Figure 1a.

We now prove that this is indeed possible without the edges crossing. Let b_1 denote the blue vertex of Δ and

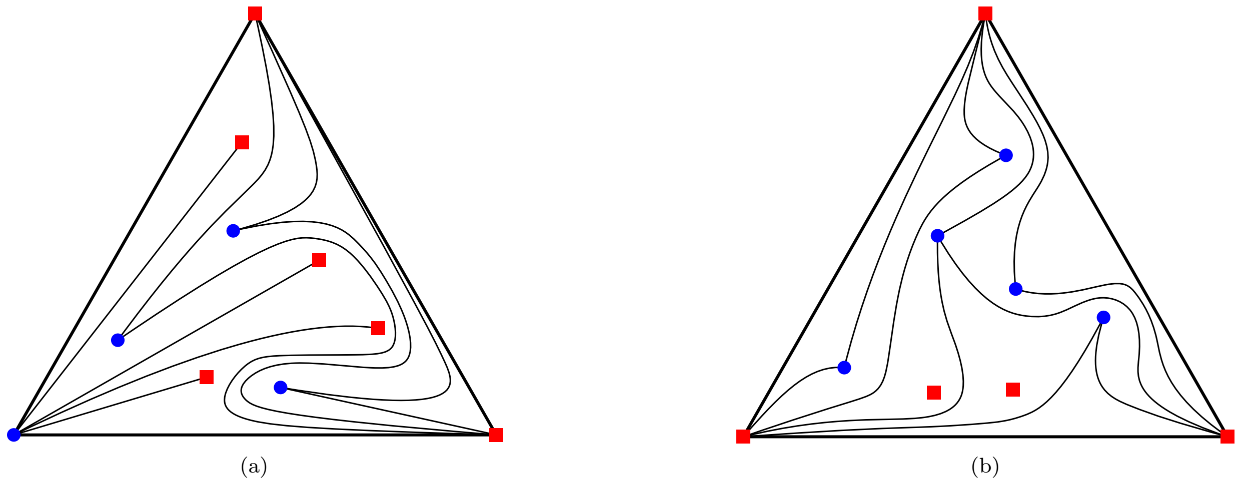

Figure 1: Part (a) illustrates Case 3.1.1 while part (b) illustrates Case 3.1.2. Projections of the red vertices of Q are marked by red squares while those of blue vertices are marked by blue disks. The vertices of Δ are also marked by red squares and blue disks.

r_1, r_2 denote the two red vertices of Δ. Let $r \in Q$ be a red vertex and p_{rb_1} be a path from $\pi(r)$ to b_1 which lies inside Δ. Clearly, $\Delta \setminus \{p_{rb_1}\}$ is path connected. Hence, after adding all edges between the projections of red points of Q and b_1, the set of remaining points of Δ is path connected. The edges between the projections of blue points of Q and the red vertices of Δ are added in order of their distance from $\overline{r_1 r_2}$, the line segment between r_1 and r_2. For each blue point b, we ensure that the region bounded by $p_{br_1} \cdot \overline{r_1 r_2} \cdot p_{br_2}^{-1}$, where $p_{br_2}^{-1}$ denotes the path that is the reverse path of p_{br_2}, does not contain the projections of any of the unprocessed blue points of Q.

Note that the case where Δ has 2 blue vertices and 1 red vertex is symmetrical to this one.

Case 3.1.2 (Δ has 3 red vertices). Consider a blue vertex b of Q. Let r_1, r_2 and r_3 denote the vertices of Δ. If there exists a a vertex v of Δ such that there exists no open halfspace of $\mathbb{R}^3$ which contains only b and v out of $\{b, r_1, r_2, r_3, o\}$; then we add edges between b and the other two vertices of Δ. If no such v exists, then we add edges between b and all three vertices of Δ. This is illustrated in Figure 1b.

We now prove that there is at most one $b \in Q$ which is connected to all three vertices of Δ. Assume that there are two blue vertices of Q, say b_1 and b_2 whose projections are joined to r_1, r_2 and r_3. Let $D = \{b_1, b_2, r_1, r_2, r_3\}$. By definition, there exists open halfspaces H_{ij}, where $i \in \{1, 2\}$ and $j \in \{1, 2, 3\}$, such that $H_{ij} \cap D = \{b_i, r_j\}$. Since b_1 and b_2 lie on the same side of Δ, there exists an open halfspace that contains only b_1 and b_2 in D. Hence, the points in D are in convex position. Since $|D| = 5$, by Theorem 1.5, there exists a Radon partition of D. Since the points of D are in convex position, there does not exist a $(1, 4)$-Radon partition of D. Thus, D has a $(2, 3)$-Radon partition. Let Y and Z be such a partition of D where $|Y| = 2$ and $|Z| = 3$. Since b_1 and b_2 lie in the interior of C, $\overline{b_1 b_2}$ does not intersect Δ. Hence, Y cannot be $\{b_1, b_2\}$. Since b_1 and b_2 lie on the same side of Δ, Y cannot be a subset of $\{r_1, r_2, r_3\}$. This proves that $Y = \{b_i, r_j\}$ for some $i \in \{1, 2\}$ and $j \in \{1, 2, 3\}$. Assume, without loss of generality, that $i = j = 1$. Then, by definition, $\overline{b_1 r_1}$ intersects the triangle described by b_2, r_2 and r_3 at some point, say p. Then, since open halfspaces are convex, any halfspace containing b_1 and r_1 contains p. Hence, it must contain at least one of b_2, r_2 or r_3. This contradicts our assumption that H_{11} exists and thus proves our claim.

Now, we prove that adding these edges is possible without them crossing. Let b^* be the blue vertex whose projection is connected to the three red vertices r_1, r_2 and r_3 of Δ. Construct paths $p_{b^*r_1}$ and $p_{b^*r_2}$ such that the region bounded by $p_{b^*r_1} \cdot r_1r_2 \cdot p_{b^*r_2}^{-1}$ contains only the projections of the blue vertices which are joined to r_1 and r_2. Now, we construct the path $p_{b^*r_3}$ such that the regions bounded by $p_{b^*r_1} \cdot r_1r_3 \cdot p_{b^*r_3}^{-1}$ and $p_{b^*r_2} \cdot r_2r_3 \cdot p_{b^*r_3}^{-1}$ contain only the projections of the blue vertices which are to be joined to r_1 and r_3 and to r_2 and r_3 respectively. We now join the projections of the other blue vertices in order of their distances from the sides of the triangle as in Case 3.1.1.

We have now constructed a planar bipartite graph $G((R,B),E)$. We complete our analysis by proving the following: for any open halfspace H of $\mathbb{R}^3$ which does not contain o and where the sets $H \cap R$ and $H \cap B$ are non-empty, there exists a $u \in H \cap \pi(R)$ and a $v \in H \cap \pi(B)$ such that $(u,v) \in E$.

If $H \cap C$ contains both red and blue vertices, then it contains a polygonal chain whose vertices are in $H \cap C$ from a red to a blue vertex of C. Hence, there exists a red-blue edge in $H \cap C$ and this edge is in E as we added all red-blue edges of C to E in the first step of construction of G. Now, without loss in generality, assume that $H \cap C \subseteq R$. Consider an open halfspace H' which is parallel to H and is contained in H. Furthermore, assume that H' contains exactly one blue point. Again, note that we can assume this since $\mathcal{H}$ is a finite set of open halfspaces. It is clear that such a halfspace indeed exists. Let b be the blue point that H' contains. Since $o \notin H$, $o \notin H'$. Hence, $\pi(b) \in H'$. If Δ is the triangle that b is mapped to, then H' contains at least one vertex of Δ.

If b was connected to exactly one vertex of Δ, then Δ has exactly one red vertex, say r_1. This implies that $r_1 \in H'$. Since $(\pi(b), r_1) \in E$, our claim is true. If b is connected to all three vertices of Δ, we are trivially done. Hence, we are left with the case where b is connected to exactly two vertices of Δ. If Δ has two red vertices and one blue vertex (as in Case 3.1.1), then our claim is true since $\pi(b)$ is connected to both the red vertices of Δ. Hence, we assume that we are in Case 3.1.2. If H' contains 2 vertices of Δ, then we are done since $\pi(b)$ is connected to at least two red vertices in this case. If not, $H' \cap (B \cup \{r_1, r_2, r_3, o\}) = \{b, r_j\}$ where r_1, r_2 and r_3 are vertices of Δ and $j \in \{1,2,3\}$. Then, by construction, it is not the case that $\pi(b)$ is connected only to the other two vertices of Δ. Hence, we infer that $(\pi(b), r_j) \in E$. This completes the proof of our claim. $\square$

By the previous theorem, we have that $\mathcal{R}(P, \mathcal{D})$ satisfies the locality condition. Hence, there exists a PTAS to solve this hitting set problem. As discussed in the beginning of this section, thi provides a PTAS to solve the hitting set problem where the ground set is $\mathbb{R}^3$ and the set of ranges is a subset of the open halfspaces of $\mathbb{R}^3$.

3.2 Dominating Terrain-Like Graphs

Definition (Terrain-Like Graph). Let $H(V,F)$ be an undirected graph where $n = |V|$. G is said to be *terrain-like* if there exists a bijective function $\pi \colon V \mapsto \{1, 2 \ldots n\}$ such that for all $\{a,b,c,d\} \subseteq V$ where $\pi(a) < \pi(b) < \pi(c) < \pi(d)$ and (a,c), (b,d) belong to the edge set, $(a,d) \in F$.

In a terrain-like graph, for $u, v \in V$, we say u *precedes* v, denoted by $u \prec v$ if $\pi(u) < \pi(v)$. We say that u *dominates* v if $u \in N_H[v]$. Here, $N_H[v]$ denotes the closed neighbourhood of the vertex v in H. It is well known that the visibility graph of terrains (which are x-monotone polygonal chains) are terrain-like [2]. Indeed, the name "terrain-like graphs" is derived from terrains.

In this part of the report, which is referenced from [1], we prove that the DOMINATING SET problem (defined below) admits a PTAS in terrain-like graphs. We prove this by showing that this problem induces a range space which satisfies the locality condition.

Definition (Dominating Set). Let $G(V, E)$ be an undirected graph and $n = |V|$. A subset I of V is called a *dominating set* if, for all $v \in V$, $N_G[v] \cap I \neq \emptyset$. Find a dominating set of smallest size of G.

Note that the DOMINATING SET problem can be restated as the HITTING SET problem: let $P = V$ and $\mathcal{D} = \{N[v] \mid v \in V\}$. Hence, finding a dominating set of a graph is equivalent to finding a hitting set of $\mathcal{R}(P, \mathcal{D})$.

Let $H(V, F)$ be a terrain-like graph. As in Section 2, let R denote an optimal hitting set of this instance and B denote the hitting set returned by a k-level local search algorithm (Algorithm 1.2). We assume, without loss of generality, that $R \cap B = \emptyset$. To use Lemma 2.5, all that we are required to do is to prove that there exists a planar bipartite graph $G((R, B), E)$ such that for all $v \in V$, there is a $r \in N_H[v] \cap R$ and $b \in N_H[v] \cap B$ such that $(r, b) \in E$. To construct such a graph, we first define a few terms and prove a simple lemma.

For a $w \in V$, if there is a $v \in R \cup B$ such that v dominates w and precedes it, then, we define $\lambda(w)$ to be the first such vertex (that is, the vertex whose $\pi(\cdot)$ is the smallest and dominates w). Similarly, if there is a $v \in R \cup B$ such that v dominates w and is preceded by w, then, we define $\rho(w)$ to be the last such vertex. Since R and B are dominating sets of H and are disjoint, for each $w \in V$, at least one of $\lambda(w)$ or $\rho(w)$ exists. Let $A_1 = \{(\lambda(w), w) \mid w \in V \text{ for which } \lambda(w) \text{ is defined}\}$. Define A_2 similarly with respect to ρ. Let $v \in V$ such that there is a $(\lambda(u), u) \in A_1$ such that $\lambda(u) \prec v \prec u$. Let $(\lambda(w), w)$ be the arc amongst these arcs for which w is the smallest. We say that $(\lambda(w), w)$ is λ-*associated* with v. Let $S_1 = \{(\lambda(w), v) \mid (\lambda(w), w) \text{ is } \lambda\text{-associated with } v\}$. We define $(w, \rho(w))$ similarly and say that $(w, \rho(w))$ is ρ-*associated* with v. Let $S_2 = \{(v, \rho(w)) \mid (w, \rho(w)) \text{ is } \rho\text{-associated with } v\}$. Finally, let $S_3 = \{(\lambda(w), \rho(w)) \mid w \in V \setminus (R \cup B) \text{ such that } \lambda(w) \text{ and } \rho(w) \text{ exists}\}$.

The following observation follows directly from the definition of terrain-like graphs.

Observation 3.2. Arcs of A_1 are non-crossing. That is, there does not exist $(\lambda(u), u), (\lambda(v), v) \in A_1$ such that $\lambda(u) \prec \lambda(v) \prec u \prec v$. Similarly, arcs of A_2 are non-crossing as well.

Hence, the graphs $G_1(V, A_1)$ and $G_2(V, A_2)$ are planar - place the vertices in V equidistantly on the x-axis and join the edges using semicircular arcs. This also implies that the graphs $\bar{G}_1(V, A_1 \cup S_1)$ and $\bar{G}_2(V, A_2 \cup S_2)$ are planar since arcs in S_1 and S_2 lie "within" arcs in A_1 and A_2 respectively. The construction of $\bar{G}_1$ is illustrated in Figure 2. On drawing the arcs in $A_1 \cup S_1$ above the x-axis and those in $A_2 \cup S_2$ below it, we get that $G_3(V, E_3)$, where $E_3 = (A_1 \cup S_1) \cup (A_2 \cup S_2)$, is planar.

Theorem 3.3. Consider the graph $G((R, B), E)$ where E be the collection of red-blue edges in $E_3 \cup S_3$. Then, G is planar. Furthermore, for all vertices $v \in V$, there is a $r \in N_H[v] \cap R$ and $b \in N_H[v] \cap B$ such

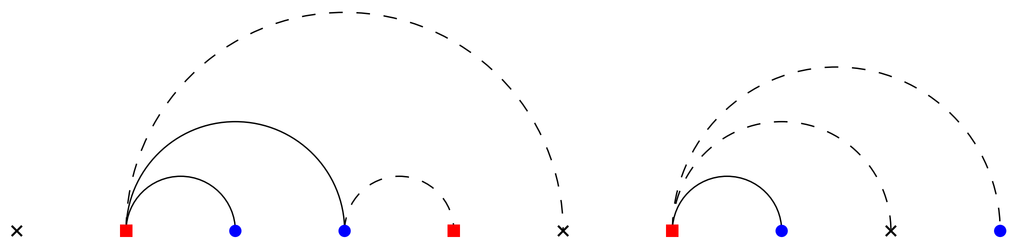

Figure 2: Vertces in R are marked by red squares, vertices in B are marked by blue discs and vertices that are in neither sets are marked by black crosses. The arcs marked by dashed lines are the ones belonging to A_1 while the ones marked by solid lines belong to S_1.

that $(r, b) \in E$.

Proof. Since G_3 is planar, it follows that $G((R, B), E)$ is planar. We now prove the second part of the claim. For any vertex $x \in R \cup B$, we let $\textsc{colour}(x)$ be *red* if $x \in R$ and *blue* if $x \in B$. Since $R \cap B = \emptyset$, this is well defined. Consider an arbitrary $v \in V$.

Case 3.3.1 ($v \notin R \cup B$). Assume that that both $\lambda(v)$ and $\rho(v)$ exist. We have two possibilities depending on the colours of these two vertices. If $\textsc{colour}(\lambda(v)) \neq \textsc{colour}(\rho(v))$, then $(\lambda(v), \rho(v)) \in E$. Since $\lambda(v)$ and $\rho(v)$ dominate v by definition, we are done. Hence, assume that $\textsc{colour}(\lambda(v)) = \textsc{colour}(\rho(v))$. Since R and B are dominating sets of V, there is a $u \in R \cup B$ with $\textsc{colour}(u) \neq \textsc{colour}(\lambda(v))$ such that u dominates v. Assume that $u \prec v$. The proof for when $v \prec u$ will follow symmetrically. Let x be the first such vertex with the above properties. Clearly, we have $\lambda(v) \prec x \prec v$. Hence, there exists an arc $(\lambda(w), w) \in A_1$ that is associated with x. If $w = v$, then, $(\lambda(w) = \lambda(v), x) \in S_1$. Since $\textsc{colour}(\lambda(v)) \neq \textsc{colour}(x)$, $(\lambda(v), x) \in E$. Furthermore, both $\lambda(v)$ and x are in $N_H[v]$. This proves our claim. If $w \neq v$, we have that $\lambda(w) \prec x \prec w \prec v$. Since $(\lambda(w), w) \in F$ and $(x, v) \in F$, $(\lambda(w), v) \in F$. Hence, $\lambda(w)$ dominates v. Since x was chosen to be the first vertex that dominated v whose colour is different from $\textsc{colour}(\lambda(v))$, $\textsc{colour}(\lambda(w)) \neq \textsc{colour}(x)$. This proves that $(\lambda(w), x) \in E$. As we have proved that both these vertices dominate v, we are done. This case is illustrated in Figure 3.

If only $\lambda(v)$ exists for v, then, there exists a u which dominates v such that $\textsc{colour}(\lambda(v)) \neq \textsc{colour}(u)$. Since $\rho(v)$ does not exist, $u \prec v$. Hence, we can proceed with the proof exactly as in the previous paragraph. Similarly, if only $\rho(v)$ exists for v, then this case will be the same as when $v \prec u$ in our discussion in the previous paragraph.

Case 3.3.2 ($v \in R \cup B$). Again, assume that both $\lambda(v)$ and $\rho(v)$ exist for v. If either of these two vertices have a different colour as compared to v, then, since they both dominate v and are neighbours of v in G, we are done. Hence, assume that $\textsc{colour}(\lambda(v)) = \textsc{colour}(v) = \textsc{colour}(\rho(v))$. Then, there exists a vertex u such that u dominates v and $\textsc{colour}(\lambda(v)) \neq \textsc{colour}(u)$. Now, we proceed exactly as in Case 3.3.1. If only one of $\lambda(v)$ and $\rho(v)$ exists for v, then, as noted in the previous case, our proof remains.

This completes our proof. $\square$

Hence, given a terrain-like graph $H(V, F)$ and a $\epsilon > 0$, by Lemma 2.5 there exists a $k \in \mathcal{O}(\epsilon^{-2})$ such that Algorithm 1.2 returns a $(1 + \epsilon)$-approximate solution for the Dominating Set problem for the instance H.

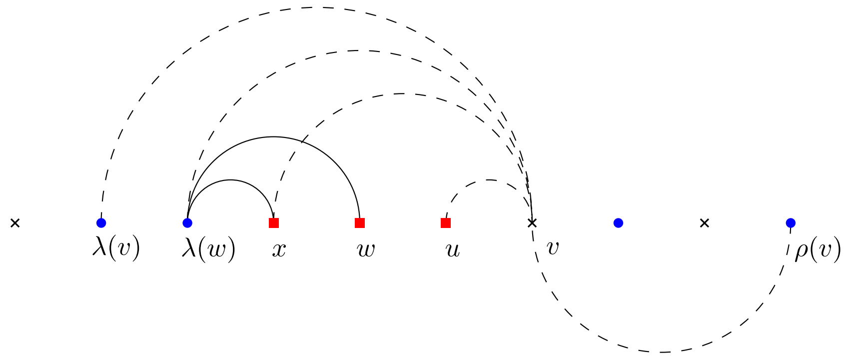

Figure 3: This illustrates Case 3.3.1. Vertces in R are marked by red squares, vertices in B are marked by blue discs and vertices that are in neither sets are marked by black crosses. The arcs marked by dashed lines are the ones belonging to F while the ones marked by solid lines belong to E.

3.3 Guarding Weakly Visible Polygons

Let $V = \{v_1, \ldots, v_n\}$ be a finite sequence of at least three points in $\mathbb{R}^2$. The *polygonal chain* defined by V is the curve specified by the line segments connecting v_i and v_{i+1} for all $1 \leq i < n$. If $v_1 = v_n$, then we say that the polygonal chain is *closed*. We define a *polygon* as the region bounded by a closed polygonal chain. A polygon is *simple* if the polygonal chain it is associated with does not cross itself. The vertices and edges of the polygonal chain corresponding to the polygon are called its vertices and edges respectively. In this section of the report, we only work with simple polygons. Furthermore, we use V to denote the vertex set of a polygon. We say that two points p, q of the polygon *see* each other if the line segment joining them lies within the polygon. For any point p in the polygon, we let $\textsc{vis}(p)$ be the set of all points seen by p. Similarly, for a subset I of the polygon, $\textsc{vis}(I)$ is the set of points of the polygon that are seen by at least one point of I.

Problem (Polygonal Vertex Guarding). Let P be a polygon. A subset I of V is *guards* V if $\textsc{vis}(I) \supseteq V$. Find a subset of V of the smallest size that guards V.

We look at a subclass of polygons, called *weakly visible* polygons for which Polygonal Vertex Guarding admits a PTAS. We do this by proving that the *visibility graph* of these types of polygons are terrain-like. From the discussion in the previous subsection, our result follows. The results presented in this subsection were obtained by Ashur *et al.* in 2019 [1].

Definition (WV Polygon). A polygon P is *weakly visible* (or WV in short) if there exists an edge (u, v) of the polygon such that for every point p on the boundary of P, there exists a q on the edge (u, v) which sees p.

Examples of WV polygons are presented in Figures 4a and 4b. We say that a WV polygon, visible from the edge (u, v), is *convex* if the internal angles measured at u and v are less than 180°. A WV polygon is *reflex* if it is not convex. We also assume, without loss in generality, that this edge lies on the x-axis.

Definition (Visibility Graph). The *visibility graph* of a polygon P is the undirected graph $G_P(V, E)$ where $E = \{(p, q) \in V^2 \mid p \text{ sees } q\}$.

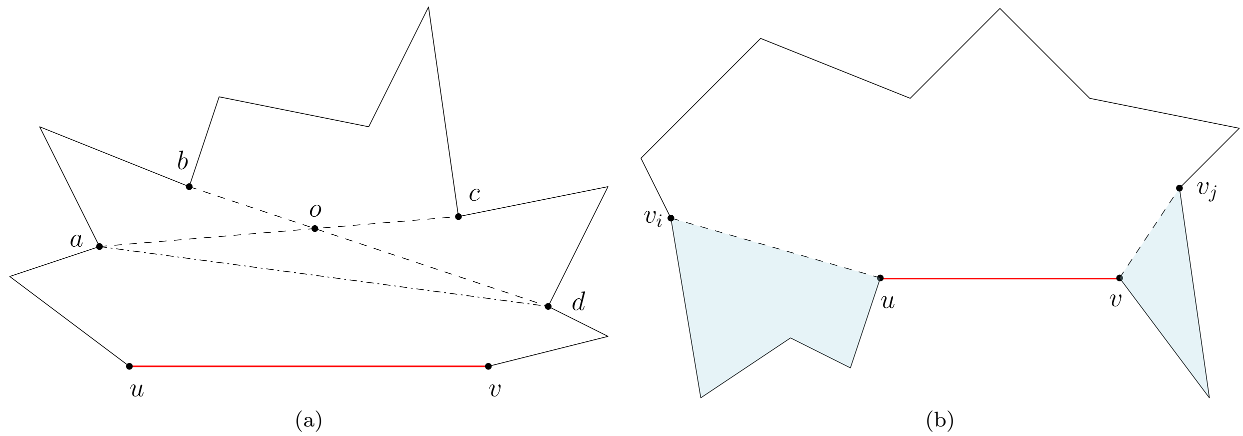

Figure 4: Part (a) illustrates convex WV polygon while part (b) illustrates a reflex WV polygon. Both these polygons are visible from the (u, v) edge which is coloured red and has a heavier stroke than the other edges. In (a), four vertices a, b, c and d are chosen such that a sees c and b sees d. The line segments between these two pairs of vertices are drawn using dashed lines. The intersection of $\overline{ac}$ and $\overline{bd}$ is marked by o. The line segment between a and d is drawn in a dashed-dotted style. In (b), v_i and v_j denote the first and last vertices excluding u and v that lie above the x-axis. The region below the x-axis upto these two vertices are coloured blue.

Observation 3.4. An instance P of the Polygonal Vertex Guarding problem is equivalent to the Dominating Set problem whose input is $G_P(V, E)$, the visibility graph of P.

We will now show that the visibility graphs of convex WV polygons are terrain-like. First, we prove a simple lemma.

Lemma 3.5. *Let P be a WV polygon that is visible from an edge (u, v). Then, no point in P lies below the x-axis.*

Proof. On the contrary, assume that there exists a vertex w of P that lies below the x-axis. Clearly, no points on the interior of the edge (u, v) can see below the x-axis. Since the interior angles at u and v are convex, they cannot see any point below the x-axis as well. Hence, w is not visible from the edge (u, v), a contradiction to our assumption that P is weakly visible from (u, v). □

Theorem 3.6. *Let P be a convex WV polygon which is visible from the edge (u, v). Order the vertices $V = \{u = v_1, v_2 \ldots v_{n-1}, v_n = v\}$ as they appear on the polygon on a clockwise traversal from u. Then, if $G_P(V, E)$ is the visibility graph of P, for all a, b, c and d in V such that $a \prec b \prec c \prec d$ where $(a, c) \in E$ and $(b, d) \in E$, $(a, d) \in E$.*

Proof. Let a, b, c and d be four such vertices of P. Since b and d lie on opposite sides of the line segment $\overline{ac}$, $\overline{bd}$ intersects $\overline{ac}$ at some point o in P. This is illustrated in Figure 4a. For two vertices v_i and v_j with $i < j$, we let $p_{v_i v_j}$ denote the path from v_i to v_j along the edges of the polygon in the clockwise direction. Consider the paths $\overline{uv}$, p_{ua}, p_{ad} and p_{dv}. Assume that $(a, d) \notin E$. Then, a and d do not see each other in P. Hence, one of these four paths must cross $\overline{ad}$. Clearly, if p_{ad} crosses $\overline{ad}$, then it must cross at least one of $\overline{ao}$ or $\overline{do}$. Since (a, c) and (b, d) are in E, this is a contradiction. By Lemma 3.5, both a and d lie above or are on the x-axis. Hence, $\overline{uv}$ does not cross the line segment $\overline{ad}$. Also, if p_{ua} crosses $\overline{ad}$, then, the point with the lowest y-coordinate that a can see lies strictly above the x-axis. Hence, it will not be visible from the (u, v) edge, a contradiction. Similarly, if p_{dv} crosses $\overline{ad}$, then d is not visible from the (u, v) edge. Hence, $(a, d) \in E$. $\square$

By the results of the previous subsection, we have the following: given a convex WV polygon P and a $\epsilon > 0$, by Lemma 2.5 there exists a $k \in \mathcal{O}(\epsilon^{-2})$ such that Algorithm 1.2 is a $(1 + \epsilon)$-approximation for the Polygonal Vertex Guarding problem for the instance P.

We now extend this result to provide a PTAS for any WV polygon. Let P be a WV polygon which is visible from its edge (u, v). As in Figure 4b, assume that P is reflex. Let v_i denote the first vertex after u which lies above the x-axis and v_j denote the last vertex before v which lies above the x-axis. We make two observations regarding the vertices that lie between u and v_i and those between v_j and v (the region marked blue in Figure 4b).

Observation 3.7. For all $v_k \in V$ such that $u \prec v_k \prec v_i$, u sees v_k. Similarly, for all $v_k \in V$ such that $v_j \prec v_k \prec v$, v sees v_k.

Observation 3.8. For all $v_k \in V$ such that $u \prec v_k \prec v_i$, $\text{VIS}(u) \supseteq \text{VIS}(v_k)$. Similarly, for all $v_k \in V$ such that $v_j \prec v_k \prec v$, $\text{VIS}(v) \supseteq \text{VIS}(v_k)$.

By Observation 3.8, we can assume that the optimal guard set does not contain any vertices between u and v_i or between v_j and v. Let P' denote the polygon described by $\{u, v_i, v_{i+1} \dots v_{j-1}, v_j, v\}$. Then, P' is a convex WV polygon. Hence, given an $\epsilon' > 0$, we can obtain a $(1 + \epsilon')$-approximate solution by running a suitable k-level local search algorithm on the $G_{P'}(V, E)$. Let I be the set obtained by adding u and v to this solution. By Observation 3.7, I guards P. Since we are adding a small number of extra guards (at most two) to our solution, given any $\epsilon > 0$, assuming that the optimal guard set is large, we can adjust ϵ' in such a way that I is a $(1 + \epsilon)$-approximate guarding set of P. We state this discussion as the final result of this section below.

Theorem 3.9. Let P be a WV polygon. Then, given any $\epsilon > 0$, there is a $k \in \mathcal{O}(\epsilon^{-2})$ such that the k-level local search algorithm run on $G_P(V, E)$ gives a $(1 + \epsilon)$-approximate guard set of P. Hence, there exists a PTAS for the Polygonal Vertex Guarding problem for WV polygons.

4 Optimality

Let $G(V, E)$ be a graph. For a $k \in \mathbb{N}$, we say that G is *k-expanding* if, for all $W \subseteq V$ where $|W| \leq k$, $|N_G(W)| \geq |W|$. In this section, we produce a family $\{G_n((R_n, B_n), E_n)\}_{n \in \mathbb{N}}$ of planar bipartite graphs which are k-expanding such that $|B_n|$ is, up to a small factor, $(1 + \frac{c}{\sqrt{k}})|R_n|$ thereby proving that Algorithm 1.2 is optimal for problems that give rise to such graphs.

A subset $\mathcal{S}$ of V is called a *separator* of G if $G[V \setminus \mathcal{S}]$ is not connected and each of its connected components have at most $\lfloor \frac{2|V|}{3} \rfloor$-many vertices. Let $\mathcal{G}$ be a family of graphs. $\mathcal{G}$ is said to have the *separator property* with parameter $s \in [0, 1]$ if there exists a $c \in \mathbb{R}$ such that for all $G(V, E) \in \mathcal{G}$ there exists a separator $\mathcal{S}$ of G with $|\mathcal{S}| \leq c \cdot |V|^{1-s}$. $\mathcal{G}$ is said to be *monotone* if, for all $G \in \mathcal{G}$ and all subgraphs G' of G, G' belongs to

$\mathcal{G}$.

Simple examples of monotone families are forests and planar graphs. We now prove that trees have the separator property with parameter $s = 1$ (and hence for all $s \in [0, 1]$). This is referenced from a proof presented on math.stackexchange by user Brandon du Preez.

Lemma 4.1. Let $T(V, E)$ be a tree. Then, there exists a separator $\mathcal{S} = \{s\}$ of T.

Proof. We define a function $C\colon V \mapsto \mathbb{N}$ by letting $C(v)$ be the maximum number of vertices a component in $T[V \setminus \{v\}]$. Let $s \in V$ be the vertex that minimizes C - that is, $C(s) \leq C(v)$ for all $v \in V$. Assume, for the sake of contradiction, that $C(s) > \lfloor \frac{2|V|}{3} \rfloor$. Then, there exists a component $T'(V', E')$ of $G[V \setminus \{s\}]$ such that $|V'| > \lfloor \frac{2|V|}{3} \rfloor$. Since T is a tree, it is connected. Hence, T' has a vertex, say w, such that $(s, w) \in E$. Moreover, since T is acyclic, this is the only such edge between vertices in V' and $V \setminus V'$. Let $T''(V'', E'') = T[V \setminus V']$. Since $|V'| > \lfloor \frac{2|V|}{3} \rfloor$,

$$|V''| < \left\lceil \frac{|V|}{3} \right\rceil \implies |V''| < |V'| \tag{4.1}$$

Now, consider $C(w)$. Every component of $T[V \setminus \{w\}]$ is a proper subgraph of T' or a subgraph of T''. Then, by Equation (4.1), we have that $C(w) < |V'|$. Hence, $C(w) < C(s)$, a contradiction to our assumption that s minimizes C. This proves that our premise that $C(s) > \lfloor \frac{2|V|}{3} \rfloor$ is incorrect. This proves our lemma since $\mathcal{S} = \{s\}$ is a separator. □

We state the celebrated *Planar Separator* Theorem [5] which proves that planar graphs have the separator property with $s = \frac{1}{2}$. While its proof uses techniques similar to the ones above, it is considerably more complicated and is thus beyond the scope of this report.

Theorem 4.2 (Planar Separator). Let $G(V, E)$ be a planar graph. Then, there exists a separator $\mathcal{S}$ of G such that $|\mathcal{S}| \leq \sqrt{8} \cdot |V|^{\frac{1}{2}}$.

In 2019, Mustafa and Jartoux [4] proved that Theorem 2.4 is optimal up to a sublinear factor. In particular, they proved the following theorem.

Theorem 4.3. There exists constants $k_0, c_0 \in \mathbb{N}$ such that for all $k \in \mathbb{N}$ where $k \geq k_0$, there exists a family of planar bipartite graphs $G_n((R_n, B_n), E_n)$ for each $n \in \mathbb{N}$ such that:

(i) G_n is k-expanding for all $n \in \mathbb{N}$.

(ii) $|R_n|, |B_n| \in \Theta(n)$ for all $n \in \mathbb{N}$.

(iii) $|B_n| \geq (1 + \frac{c_0}{\sqrt{k}})|R_n| - o(n)$ as $n \to \infty$.

In fact, they prove a much more general statement:

Theorem 4.4. Given a positive integer d, there exists constants $k_d, c_d \in \mathbb{N}$ such that for all $k \in \mathbb{N}$ where $k \geq k_d$, there exists a family of graphs $G_n((R_n, B_n), E_n)$ for each $n \in \mathbb{N}$ such that:

(i) G_n, and all its subgraphs, have the separator property for $s = \frac{1}{d}$ for all $n \in \mathbb{N}$.

(ii) G_n is k-expanding for all $n \in \mathbb{N}$.

(iii) $|R_n|, |B_n| \in \Theta(n)$ for all $n \in \mathbb{N}$.

(iv) $|B_n| \geq (1 + c_d \cdot k^{-\frac{1}{d}})|R_n| - o(n)$ as $n \to \infty$.

Since planar graphs have the separator property when $s = \frac{1}{2}$, Theorem 4.3 is a special case of Theorem 4.4 (when $d = 2$). The proof for arbitrarily large d is a little more unwieldy, but follows the same structure as our proof.

We now construct a family of planar bipartite graphs $\{G_n((R_n, B_n), E_n)\}_{n\in\mathbb{N}}$ which have the three properties listed above. For each $n \in \mathbb{N}$, we call R_n to be the set of *red* vertices and B_n to be the set of *blue* vertices. Since this construction is geometric (each vertex of G_n is a point in $\mathbb{R}^2$), we often interchange the terms "vertices" and "points". For a point $p \in \mathbb{R}^2$, we let $x(p)$ denote its x-coordinate and $y(p)$ denote its y-coordinate.

Let $L \geq 2$ be a natural number. We will, at the end of this section, fix it to be a function of k. We construct G_n in two steps. First, fix a $\overrightarrow{\tau} \in \mathbb{R}^2$ with $x(\overrightarrow{\tau}) = y(\overrightarrow{\tau})$. Consider Ξ, the $L \times L$ regular integer grid in $\mathbb{R}^2$ which is anchored at $\overrightarrow{\tau}$. This has L^2-many *cells* and $(L+1)^2$-many points. The bottom left vertex of a cell C, is called its *anchor* vertex and is denoted by $\textsc{anc}(C)$. We use $x(C)$ and $y(C)$ to denote the x-coordinate and y-coordinate of $\textsc{anc}(C)$ respectively. Clearly, each vertex of Ξ, apart from those with one of the coordinate values equal to $L + x(\overrightarrow{\tau})$, is an anchor to exactly one cell. This cell is called its *top* cell. The cell whose anchor vertex is $\overrightarrow{\tau}$ is called the *lowest* cell of Ξ while the red vertex at τ is called the *lowest* vertex of Ξ. A *block slab* of Ξ is a subset of its cells which are in the same row or column.

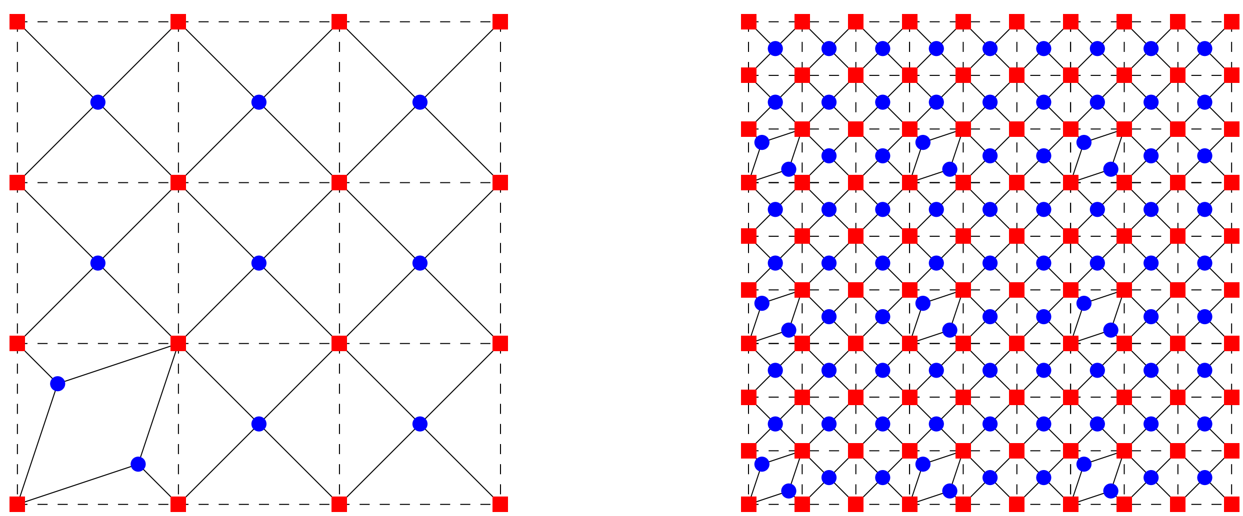

(a) (b)

Figure 5: Part (a) illustrates the construction of $G^{\overrightarrow{\tau}}$ with $L = 3$ while part (b) illustrates the construction of G_n with $t = 3$ and $G^{\overrightarrow{\tau}}$ as translates. Vertices of $R^{\overrightarrow{\tau}}$ and R_n are marked by red squares while those of $B^{\overrightarrow{\tau}}$ and B_n are marked by blue disks. The outline of the grids are marked using dashed lines while edges of $G^{\overrightarrow{\tau}}$ and G_n are marked using solid lines.

We now define a planar bipartite graph $G^{\overrightarrow{\tau}}((R^{\overrightarrow{\tau}}, B^{\overrightarrow{\tau}}), E_{\overrightarrow{\tau}})$ on Ξ (Figure 5a). $R^{\overrightarrow{\tau}}$ is just the set of points which define Ξ. We place a blue vertex at the center of each cell of Ξ except for the lowest cell. In the lowest cell, we place blue vertices at $(x(\overrightarrow{\tau}) + \frac{1}{4}, y(\overrightarrow{\tau}) + \frac{3}{4})$ and $(x(\overrightarrow{\tau}) + \frac{3}{4}, y(\overrightarrow{\tau}) + \frac{1}{4})$. Construct an edge between the blue vertex in a cell (other than the lowest cell) to the four red vertices that define it. For all such vertices, we refer to the anchor vertex of the cell that contains it as its *anchor* vertex. Join the blue vertex at $(x(\overrightarrow{\tau}) + \frac{1}{4}, y(\overrightarrow{\tau}) + \frac{3}{4})$ to all the vertices that define the lowest cell except for $(x(\overrightarrow{\tau}) + 1, y(\overrightarrow{\tau}))$ and the blue vertex at $(x(\overrightarrow{\tau}) + \frac{3}{4}, y(\overrightarrow{\tau}) + \frac{1}{4})$ to all the vertices that define the lowest cell except for $(x(\overrightarrow{\tau}), y(\overrightarrow{\tau}) + 1)$. It is clear that the graph so constructed is planar.

We now complete the construction of $G_n((R_n, B_n), E_n)$. Let $t \in \mathbb{N}$. We will, along with L, fix it to be a function of k and n at the end of this section. $G_n((R_n, B_n), E_n)$ consists of a $t \times t$ grid containing t^2-many *translates* of $G^{\overrightarrow{\tau}}$. Each translate is indexed and anchored by $\overrightarrow{\tau} \in \{0, 1 \dots t-1\}^2$. R_n (B_n) is just the union of all red (blue) vertices in these translates. Note that we identify the red vertices that share the same coordinates in these translates as one. E_n too is just the union of the edges defined in $G^{\overrightarrow{\tau}}$. The author refers the reader to Figure 5b for an illustration of this construction.

The subset of vertices in R_n whose x-coordinate or y-coordinate is equal to tL is called R_b. Let S be a subset of cells in the grid underlying G_n. The *bounding box* of S is defined to be the smallest rectangular subgrid that contains S. We make the following observations regarding the size of R_n and B_n.

$$|R_n| = (tL + 1)^2 \tag{4.2}$$

$$|B_n| = t^2(L^2 + 1) \tag{4.3}$$

We show the following: for all $B' \subseteq B_n$ where $|B'|$ is bounded above by some function of L (which we will manipulate later to ensure that the function itself is at most k), $|N_{G_n}(B')| \geq |B'|$.

Fix a $B' \subseteq B_n$ and let $R' = N_{G_n}(B')$. Cells of G_n which contain vertices from B' are called *non-empty* while those that do not contain such vertices are said to be *empty*. A vertex of R' which also belongs to R_b or whose top cell is empty is called a *boundary* vertex. The total number of boundary vertices in G_n is denoted by δ. For a translate $G^{\overrightarrow{\tau}}$, we define $d_{\overrightarrow{\tau}} = 1$ if both the blue vertices in the lowest cell of $G^{\overrightarrow{\tau}}$ is in B' and define $d_{\overrightarrow{\tau}} = 0$ otherwise. Moreover, we let $\delta_{\overrightarrow{\tau}}$ to be the number of boundary vertices that are in $G^{\overrightarrow{\tau}}$.

We prove a small lemma comparing δ to the size of B'.

Lemma 4.5. $\delta \geq \sqrt{\frac{|B'|}{2}}$.

Proof. Let S be the set of non-empty cells of G and $\mathcal{B}$ be its bounding box. Let the length and breadth of $\mathcal{B}$ be p and q respectively. Clearly, each slab of $\mathcal{B}$ contains at least 1 boundary vertex. Hence, $\delta \geq p$ and $\delta \geq q$. This implies that

$$\delta^2 \geq pq \geq |S| \tag{4.4}$$

Moreover, since each cell in S contains at most 2 blue vertices, $2|S| \geq |B'|$. Plugging this inequality into Equation (4.4), we have

$$\delta^2 \geq \frac{|B'|}{2} \implies \delta \geq \sqrt{\frac{|B'|}{2}}$$

This proves our lemma. ∎

For each vertex in B', its anchor vertex belongs to R'. We now prove that there is an injective mapping from B' to R' if $|B'|$ is "large enough". In each non-empty cell, map a blue vertex to its anchor vertex. For all $G^{\overrightarrow{\tau}}$ where $d_{\overrightarrow{\tau}} = 1$, exactly one blue vertex remains unmapped. Meanwhile, the boundary vertices of G_n remain unmapped to a blue vertex. For those $G^{\overrightarrow{\tau}}$ which have $\delta_{\overrightarrow{\tau}} \geq 2$, map one of them to the unmapped blue vertex of its lowest cell. Let δ' be the number of boundary vertices are still unmapped to a blue vertex. Clearly, $\delta' \geq \frac{\delta}{2}$. By Lemma 4.5, we have

$$\delta' \geq \frac{1}{2} \cdot \sqrt{\frac{|B'|}{2}} \tag{4.5}$$

We are now left with those $G^{\overrightarrow{\tau}}$ where $d_{\overrightarrow{\tau}} = 1$ and $\delta_{\overrightarrow{\tau}} \leq 1$. We prove that each such translate must contain at least $\frac{L^2}{2}$-many vertices from B'.

Lemma 4.6. Let $G^{\overrightarrow{\tau}}$ be a translate such that $\delta_{\overrightarrow{\tau}} \leq 1$ and $d_{\overrightarrow{\tau}} = 1$. Then, $|B' \cap B^{\overrightarrow{\tau}}| \geq \frac{L^2}{2}$.

Proof. First, assume that $\delta_{\overrightarrow{\tau}} = 0$. Hence, $G^{\overrightarrow{\tau}}$ does not contain any boundary vertices. Clearly, this implies that every vertex in $B^{\overrightarrow{\tau}}$ is in B' proving that $|B'| \geq L^2 \geq \frac{L^2}{2}$.

Now, assume that $\delta_{\overrightarrow{\tau}} = 1$. Then, $G^{\overrightarrow{\tau}}$ has exactly one boundary vertex, say v_r. Since $d_{\overrightarrow{\tau}} = 1$, B' contains both of the blue vertices of the lowest cell of $G^{\overrightarrow{\tau}}$. Assume, without loss in generality, that the lowest vertex of $G^{\overrightarrow{\tau}}$ is placed at the origin. Hence, v_r is not the lowest vertex of $G^{\overrightarrow{\tau}}$ implying that at least one of $x(v_r)$ or $y(v_r)$ are positive. We assume that $y(v_r) > 0$. The proof for the other case will follow symmetrically. Consider the block slab Ξ' containing cells whose anchor vertices lie on the x-axis.

Claim 4.6.1. Every cell of Ξ' is non-empty.

Proof of Claim. Assume, for the sake of contradiction, that there exists a cell in Ξ' which is empty. Let C be such a cell with the lowest x-coordinate. Since $d_{\overrightarrow{\tau}} = 1$, C is not the lowest cell of $G^{\overrightarrow{\tau}}$. Hence, there exists a cell $C' \in \Xi'$ such that $x(C') = x(C) - 1$. Moreover, by our assumption, C' is non-empty. This implies that $\textsc{anc}(C) \in R'$. Since $y(C) = 0$, $\textsc{anc}(C) \neq v_r$. This contradicts our assumption that v_r is the only boundary vertex of $G^{\overrightarrow{\tau}}$ and completes the proof of this claim. $\Diamond$

We now prove that "most" cells of $G^{\overrightarrow{\tau}}$ are non-empty.

Claim 4.6.2. Every cell in $G^{\overrightarrow{\tau}}$ whose x-coordinate is different from that of v_r is non-empty.

Proof of Claim. Consider a cell $C \in G^{\overrightarrow{\tau}}$ such that $x(C) \neq x(y_r)$. Consider the block slab Ξ'' which contains the cells in $G^{\overrightarrow{\tau}}$ with x-coordinates same as that of C. Since the cell in Ξ'' whose y-coordinate is 0 is non-empty by Claim 4.6.1, using arguments similar to the ones presented in Claim 4.6.1, we can prove that every cell in Ξ'' is non-empty. In particular, C is non-empty. This proves our claim. $\Diamond$

By Claim 4.6.2, we have that the $L(L-1)$-many cells whose x-coordinates are different from that of v_r are non-empty. Hence,

$$|B' \cap B^{\overrightarrow{\tau}}| \geq L(L-1) \geq \frac{L^2}{2}$$

This completes the proof of the lemma. $\square$

Since those $G^{\overrightarrow{\tau}}$ where $d_{\overrightarrow{\tau}} = 1$ and $\delta_{\overrightarrow{\tau}} \leq 1$ have exactly one unmapped blue vertex, by Lemma 4.6, there are at most $\frac{2|B'|}{L^2}$-many unmapped blue vertices left in G_n. Now, note that

$$|B'| \leq 2^{-5}L^4 \implies \frac{4|B'|^2}{L^4} \leq \frac{1}{4} \cdot \frac{|B|'}{2} \implies \frac{2|B'|}{L^2} \leq \frac{1}{2} \cdot \sqrt{\frac{|B|'}{2}} \implies \frac{2|B'|}{L^2} \leq \delta' \tag{4.6}$$

That is, Equation (4.6) implies that if $|B'| \leq 2^{-5}L^4$, then the number of unmapped blue vertices is at most the number of unmapped red vertices.

To complete the proof of Theorem 4.3, we must set appropriate values for L and t given a k and n. Consider:

(i) $L = \lceil (2^{-5}k)^{\frac{1}{4}} \rceil$.

(ii) $t = \lceil \frac{\sqrt{n}}{L} \rceil$.

Then, by Equations (4.2) and (4.2), we have,

$$|R_n| = (tL+1)^2 = n + o(n) \tag{4.7}$$

$$|B_n| = t^2(L^2+1) = n + o(n) \tag{4.8}$$

Moreover,

$$\begin{aligned}
\frac{|B_n|}{|R_n|} &= \frac{t^2(L^2+1)}{(tL+1)^2} \\
&\geq \frac{n+t^2}{(tL+1)^2} && (t \geq \frac{\sqrt{n}}{L}) \\
&\geq \frac{n+t^2}{(\sqrt{n}+L+1)^2} && (t \leq \frac{\sqrt{n}}{L}+1) \\
&= \frac{n+t^2}{n+\sqrt{n}} - o(1) \text{ as } n \to \infty \\
&= 1 + \frac{n}{L^2(n+\sqrt{n})} - o(1) \text{ as } n \to \infty \\
&= 1 + \frac{1}{L^2} - o(1) \text{ as } n \to \infty \\
&\geq 1 + \frac{c_0}{\sqrt{k}} - o(1) \text{ as } n \to \infty && (4.9)
\end{aligned}$$

Here, $c_0 = 4\sqrt{2}$. By Equations (4.7) and (4.9), we have have that

$$|B_n| \geq (1 + \frac{c_0}{\sqrt{k}})|R_n| - o(n) \text{ as } n \to \infty \tag{4.10}$$

Equations (4.7), (4.8) and (4.10) were exactly the properties of planar bipartite graphs that we were after for the proof of Theorem 4.3. This completes this section and thus the report.

Acknowledgments

The author would like to thank Sutanu Roy (SMS, NISER) for the opportunity to work on this topic as part of his 8^{th} and 9^{th} semester projects. He would also like to extend his thanks to Aritra Banik (SCPS, NISER) and Rajiv Raman (LIMOS, UCA) for their helpful discussions on the topics presented in Subsection 3.1.

References

[1] Stav Ashur et al. "Terrain-Like Graphs: PTASs for Guarding Weakly-Visible Polygons and Terrains". In: *Approximation and Online Algorithms - 17th International Workshop, WAOA 2019, Munich, Germany, September 12-13, 2019, Revised Selected Papers*. Ed. by Evripidis Bampis and Nicole Megow. Vol. 11926. Lecture Notes in Computer Science. Springer, 2019, pp. 1–17. DOI: 10.1007/978-3-030-39479-0_1.

[2] Boaz Ben-Moshe, Matthew J. Katz, and Joseph S. B. Mitchell. "A Constant-Factor Approximation Algorithm for Optimal 1.5D Terrain Guarding". In: *SIAM J. Comput.* 36.6 (2007), pp. 1631–1647. DOI: 10.1137/S0097539704446384.

[3] Greg N. Frederickson. "Fast Algorithms for Shortest Paths in Planar Graphs, with Applications". In: *SIAM J. Comput.* 16.6 (1987), pp. 1004–1022. DOI: 10.1137/0216064.

[4] Bruno Jartoux and Nabil H. Mustafa. "Optimality of Geometric Local Search". In: *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*. Ed. by Bettina Speckmann and Csaba D. Tóth. Vol. 99. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, 48:1–48:15. DOI: 10.4230/LIPIcs.SoCG.2018.48. URL: https://doi.org/10.4230/LIPIcs.SoCG.2018.48.

[5] Richard J. Lipton and Robert Endre Tarjan. "Applications of a Planar Separator Theorem". In: *SIAM J. Comput.* 9.3 (1980), pp. 615–627. DOI: 10.1137/0209046. URL: https://doi.org/10.1137/0209046.

[6] Nabil H. Mustafa. *Lecture 2: Radon's Theorem*. URL: https://perso.esiee.fr/~mustafan/TechnicalWritings/math-lec2.pdf.

[7] Nabil H. Mustafa and Saurabh Ray. "Improved Results on Geometric Hitting Set Problems". In: *Discret. Comput. Geom.* 44.4 (2010), pp. 883–895. DOI: 10.1007/s00454-010-9285-9.