Some Problems Related to Good Illumination*

Manuel Abellanas[1]**, Antonio Bajuelos[2]***, Inês Matos[2]***

[1] Facultad de Informática, Universidad Politécnica de Madrid, Spain
mabellanas@fi.upm.es
[2] Departamento de Matemática & CEOC, Universidade de Aveiro, Portugal
{leslie,ipmatos}@mat.ua.pt

Abstract. A point p is 1-well illuminated by a set of n point lights if there is, at least, one light interior to each half-plane with p on its border. We consider the illumination range of the lights as a parameter to be optimized. So we minimize the lights' illumination range to 1-well illuminate a given point p. We also present two generalizations of 1-good illumination: the orthogonal good illumination and the good Θ-illumination. For the first, we propose an optimal linear time algorithm to optimize the lights' illumination range to orthogonally well illuminate a point. We present the E-Voronoi Diagram for this variant and an algorithm to compute it that runs in $\mathcal{O}(n^4)$ time. For the second and given a fixed angle $\Theta \leq \pi$, we present a linear time algorithm to minimize the lights' illumination range to well Θ-illuminate a point.

Key words: Computational Geometry, Limited Range Illumination, Good Illumination, E-Voronoi Diagrams

1 Introduction and Related Works

Visibility and illumination have been a main topic for different papers in the area of Computational Geometry (for more information on the subject, see Asano et. al [4] and Urrutia [16]). However, most of these problems deal with ideal concepts. For instance, light sources have some restrictions as they cannot illuminate an infinite region since their light naturally fades as the distance grows. This is also the case of cameras and robot vision systems, both have severe visibility range restrictions since they cannot observe with sufficient detail far away objects. We present some of these illumination problems adding several restrictions to make them more realistic. Each light source has limited illumination range so that their illuminated regions are delimited. We use a definition of limited visibility due to Ntafos [14] as well as a concept related to this type of problems,

* When this paper was finished, the third author was supported by a FCT fellowship, grant SFRH/BD/28652/2006

** Supported by grant TIC2003-08933-C02-01, MEL-HP2005-0137 and partially supported by CAM:P-DPI-000235-0505

*** Supported by CEOC through *Programa* POCTI, FCT, co-financed by EC fund FEDER and by Acção Integrada Luso-Espanhola No. E-77/06

the t-good illumination due to Canales et. al [3,7]. This study is solely focused on an optimization problem related to limited range illumination. In its original definition [1], a point is 1-well illuminated if it lies in the interior of the convex hull of a set of light sources.

This paper is structured as follows. In the next section we formalize the 1-good illumination and propose an algorithm to calculate the Minimum Embracing Range (MER) of a point in the plane. Sections 3 and 4 are devoted to extensions of 1-good illumination. In section 3 we present the orthogonal good illumination and propose an algorithm to compute the MER to orthogonally well illuminate a point. We follow presenting the E-Voronoi Diagram for this variant and an algorithm to compute it. In section 4 we extend 1-good illumination to cones and make a brief relation between this variant and the Maxima Problem [5,13]. We conclude this paper in section 5.

1.1 Preliminaries and Problem Definition

Let $F = \{f_1, f_2, \ldots, f_n\}$ be a set of light sources in the plane that we call sites. Each light source $f_i \in F$ has limited illumination range $r > 0$, so f_i only illuminates objects that are within the circle centered at f_i with radius r. The next definitions follow from the notation introduced by Chiu and Molchanov [9]. The set $\mathrm{CH}(F)$ represents the convex hull of the set F.

Definition 1. *A set of light sources F is called an embracing set for a point p in the plane if p lies in the interior of the* $\mathrm{CH}(F)$*.*

Definition 2. *A site $f_i \in F$ is an embracing site for a point p if p lies in the interior of the convex hull formed by f_i and by all the sites of F closer to p than f_i.*

As there may be more than one embracing site per point, our main goal is to compute a Closest Embracing Site for a given point p since we are trying to minimize the light sources' illumination range (see Fig. 1(a) and Fig. 1(b)).

Definition 3. *Let F be a set of n light sources. A set formed by a closest embracing site for p, f_i, and all the lights sources closer to p than f_i is called a minimal embracing set for p.*

Definition 4 ([7]). *Let F be a set of n light sources. We say that a point p in the plane is t-well illuminated by F if every open half-plane with p on its border contains at least t light sources of F illuminating p.*

This definition tests the light sources' distribution in the plane so that the greater the number of light sources in every open half-plane containing the point p, the better the illumination of p. This concept can also be found under the name of $\triangle$-guarding [15] or well-covering [10]. The motivation behind this definition is the fact that, in some applications, it is not sufficient to have one point illuminated but also some of its neighbourhood [10].

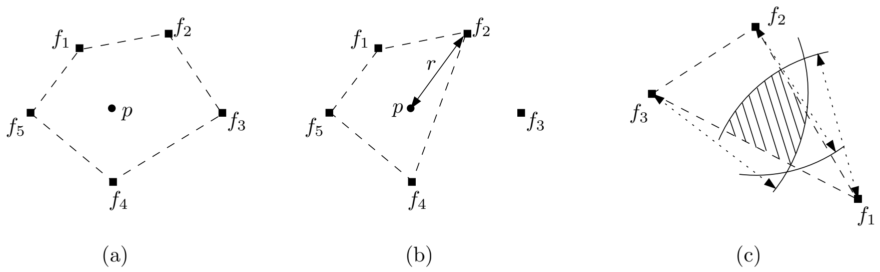

Fig. 1. (a) The light sources f_2 and f_3 are embracing sites for point p. (b) The light source f_2 is the closest embracing site for p and its illumination range is $r = d(p, f_2)$. The set $\{f_1, f_2, f_4, f_5\}$ is a minimal embracing set for p. (c) $A_r^E(f_1, f_2, f_3)$ is the shaded open area, so every point that lies inside it is 1-well illuminated by f_1, f_2 and f_3.

Let $C(f_i, r)$ be the circle centered at f_i with radius r and let $A_r(f_i, f_j, f_k)$ denote the r-illuminated area by the light sources f_i, f_j and f_k. It is easy to see that $A_r(f_i, f_j, f_k) = C(f_i, r) \cap C(f_j, r) \cap C(f_k, r)$. We use $A_r^E(f_i, f_j, f_k) = A_r(f_i, f_j, f_k) \cap \text{int}(\text{CH}(f_i, f_j, f_k))$ to denote the illuminated area embraced by the light sources f_i, f_j and f_k.

Definition 5. *Let F be a set of light sources, we say that a point p is* 1*-well illuminated if there exists a set of three light sources $\{f_i, f_j, f_k\} \in F$ such that $p \in A_r^E(f_i, f_j, f_k)$ for some range $r > 0$.*

Definition 6. *Given a set F of n light sources, we call Minimum Embracing Range to the minimum range needed to* 1*-well illuminate a point p or a set of points S in the plane, respectively* $\text{MER}(F, p)$ *or* $\text{MER}(F, S)$.

Fig. 1(c) illustrates Definition 5. Since the set F is clear from the context, we will use "MER of p" instead of $\text{MER}(F, p)$ and "MER of S" instead of $\text{MER}(F, S)$. Once we have found the closest embracing site for a point p, its MER is given by the euclidean distance between the point and its closest embracing site. Computing the MER of a given point p is important to us. The minimum illumination range that the light sources of the minimal embracing set need to 1-well illuminate p is its MER.

As an example of application, suppose that a user needs to be covered by at least one transmitter in every half-plane that passes through him in order to be well located. In such situation, the user is 1-well illuminated by the transmitters. Suppose now that we have a group of users moving from time to time, while someone has to adapt the transmitters' power so that the users don't get lost. The power of all the transmitters is controlled by a gadget that allows a constant change of the power of all the transmitters at once. But the more power, the more expensive the system is. So, it is required to know which is the minimum power that 1-well illuminates all users every time they move, that is, this problem is solved by computing the MER of each user.

2 1-Good Illumination

Let F be a set of n light sources in the plane and p a point we want to 1-well illuminate.

Definition 7. *We call Closest Embracing Triangle for a point p, $\mathrm{CET}(p)$, to a set of three light sources of F containing p in the interior of the triangle they define, such that one of these light sources is a closest embracing site for p and the other two are closer to p than its closest embracing site.*

The objective of this section is to compute the value of the MER of p and a $\mathrm{CET}(p)$. The Nearest Neighbourhood Embracing Graph (NNE-graph) [9] consists of a set of vertices V of the graph where each vertex $v \in V$ is connected to its first nearest neighbour, its second nearest neighbour, ..., until v is an interior point to the convex hull of its nearest neighbours. Chan et al. [8] present several algorithms to construct the NNE-graph. A closest embracing site for p can be obtained in linear time using this graph. The algorithm we present in this section has the same time complexity but it has the advantage of also computing a $\mathrm{CET}(p)$.

2.1 Minimum Embracing Range of a 1-well Illuminated Point

To compute the MER of p, we start by computing the distances from p to all the light sources. Afterwards, we compute the median of all the distances in linear time [6]. Depending on this value, we split the light sources in two halves: the set F_c that contains the closest half to p and the set F_f that contains the furthest half. We check whether $p \in \mathrm{int}(\mathrm{CH}(F_c))$, what is equivalent to test if F_c is an embracing set for p (see Fig. 2(a)). If the answer is negative, we recurse adding the closest half of F_f. Otherwise (if $p \in \mathrm{int}(\mathrm{CH}(F_c))$), we recurse halving F_c (see Fig. 2(b)). This logarithmic search runs until we find the light source $f_p \in F$ and the subset $F^E \subseteq F$ such that $p \in \mathrm{int}(\mathrm{CH}(F^E))$ but $p \notin \mathrm{int}(\mathrm{CH}(F^E \setminus \{f_p\})$. The light source f_p is the closest embracing site for p and its MER is $r = d(f_p, p)$ (see Fig. 2(c)).

On each recursion, we have to check whether $p \in \mathrm{int}(\mathrm{CH}(F')), F' \subseteq F$. This can be done in linear time [12] if we choose the set of points carefully so that each point is studied only once. When we have the closest embracing site for p, f_p, we find two other vertices of a $\mathrm{CET}(p)$ in linear time as follows. Consider the circle centered at p of radius r and the line $\overline{pf_p}$ that splits the light sources inside the circle in two sets. Note that if f_p is the closest embracing site for p then there is an empty semicirle. A $\mathrm{CET}(p)$ has f_p and two other light sources in the circle as vertices. Actually, any pair of light sources f_l, f_r interior to the circle such that each lies on a different side of the line $\overline{pf_p}$ verifies that $p \in \mathrm{int}(\mathrm{CH}(f_l, f_p, f_r))$.

Proposition 1. *Given a set F of n light sources and a point p in the plane, the algorithm just presented computes the* MER *of p and a Closest Embracing Triangle for it in $\Theta(n)$ time.*

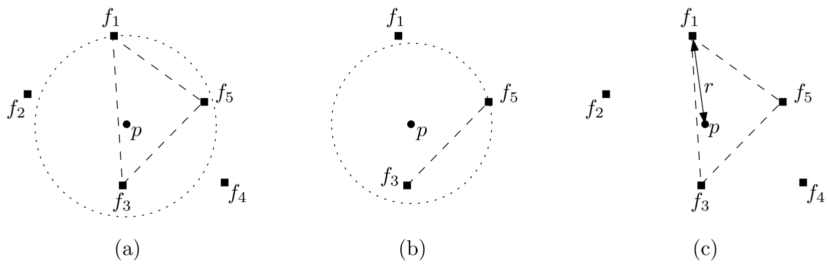

Fig. 2. (a) Point $p \in \text{int}(\text{CH}(F_c))$, where $F_c = \{f_1, f_3, f_5\}$ and $F_f = \{f_2, f_4\}$. (b) Point $p \notin \text{int}(\text{CH}(F_c))$, where $F_c = \{f_3, f_5\}$ and $F_f = \{f_1\}$. (c) The set $\{f_1, f_3, f_5\}$ is a minimal embracing set for p and the MER of p is r.

Proof. Let F be a set of n light sources. The distances from p to all the light sources can be computed in linear time. Computing the median also takes linear time [6], as well as splitting F in two halves. Checking if $p \in \text{int}(\text{CH}(F'))$, $F' \subseteq F$, is linear on the number of light sources in F'. So the total time for this logarithmic search is $\mathcal{O}(n + \frac{n}{2} + \frac{n}{4} + \frac{n}{8} + ...) = \mathcal{O}(n)$. Therefore, we find the closest embracing site for p in linear time. So this algorithm computes the MER of p and a $\text{CET}(p)$ in total $\mathcal{O}(n)$ time.

All the light sources of F must be analyzed at least once since they are all candidates to be the closest embracing site for a point p. Knowing this, we have $\Omega(n)$ as a lower bound which makes the linear complexity of this algorithm optimal. □

The decision problem is trivial after the MER of p is computed. Point p is 1-well illuminated if the given illumination range is greater or equal to its MER.

3 Orthogonal Good Illumination

This section is devoted to a variant of the 1-good illumination of minimum range using quadrants, the orthogonal good illumination. We propose an optimal linear time algorithm to compute the MER of an orthogonally well illuminated point, as well as a minimal embracing set for it. Next we present the E-Voronoi Diagram [2] for this variant, as well as an algorithm to compute it that runs in $\mathcal{O}(n^4)$ time.

An oriented quadrant is defined by two orthogonal rays that are axis-parallel. The next definition is illustrated in Fig. 3(a).

Definition 8. *Let F be a set of n light sources in the plane. We say that a point p in the plane is orthogonally well illuminated if there is, at least, one light source interior to each of the four oriented quadrants with origin at p,* NE, NW, SW *and* SE.

As it is clear from the context of this section, orthogonal good illumination will be referred to just as good illumination. The main structure in this section

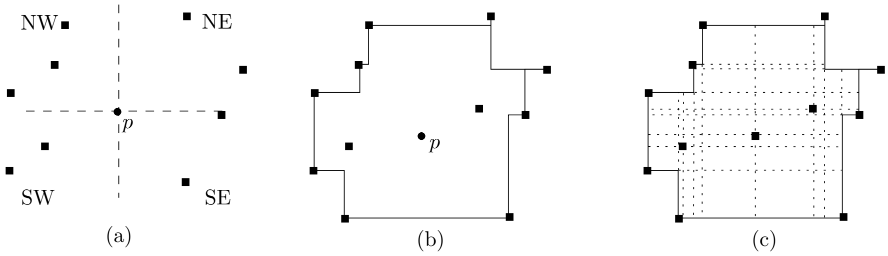

Fig. 3. (a) Point p is orthogonally well illuminated because all oriented quadrants centered at p are non-empty. (b) Point p is interior to the orthogonal convex hull of F, so it is orthogonally well illuminated. (c) An orthogonal convex hull decomposed into several rectangles.

is the orthogonal convex hull (see Karlsson and M. Overmars [11]). The convex hull of a set of points is the smallest convex region that contains it. The prefix orthogonal means that the convexity is defined by axis-parallel point connections. When $|F| \geq 4$ there is, at least, one light source of F in each quadrant centered at a point interior to the orthogonal convex hull of F. So the interior points to the orthogonal convex hull of F are well illuminated (see Fig. 3(b)).

3.1 Minimum Embracing Range of an Orthogonally Well Illuminated Point

Let F be a set of n light sources and p a point we want to well illuminate. The decision problem is easy to solve, we have to check if there is, at least, one light source interior to each of the four quadrants centered at p. If there is an empty quadrant then p is not well illuminated. Since there must be a light source in each quadrant centered at p, a minimal embracing set for p has four light sources. Let us consider the closest light source to p in each quadrant, the closest embracing site for p is the furthest of these four. The MER of p is given by the distance between p and its closest embracing site.

Proposition 2. *Given a set F of n light sources and a point p in the plane, computing a minimal embracing set for p and its* MER *takes $\Theta(n)$ time.*

Proof. Given a set F of n light sources and a point p in the plane, checking if all quadrants are empty can be done while searching for the closest light source to point p in each quadrant. This search is obviously linear on the number of light sources, while computing the MER is constant. So the total time for computing a minimal embracing set for p and its MER is $\mathcal{O}(n)$.

Since all the light sources of F are candidates to be the closest embracing site for a point p in the plane, we have to search through them all. Knowing this, we have $\Omega(n)$ as a lower bound which makes the linear complexity of this algorithm optimal. □

3.2 The E-Voronoi Diagram

When studying problems related to good illumination, one question naturally pops up: how do we preprocess the set F so that it is straightforward to know which is the closest embracing site for each point in the plane? Having such a structure would be of a great help to efficiently answer future queries. This problem is already solved by Abellanas et al. [2] when considering the usual 1-good illumination.

Definition 9 ([2]). *Let F be a set of n light sources in the plane. For every light source $f_i \in F$, the* E-*Voronoi region of f_i with respect to the set F is the set* $\text{E-VR}(f_i, F) = \{x \in \mathbb{R}^2 : f_i$ *is the closest embracing site for* $x\}$.

The region $\text{E-VR}(f_i, F)$ will be denoted by $\text{E-VR}(f_i)$ since the set F is clear from the context. The union of all the E-Voronoi regions ($\bigcup_{f_i \in F} \text{E-VR}(f_i)$) is called the E-Voronoi Diagram of F. So if $p \in \text{E-VR}(f_i)$ then the MER of p is the distance between f_i and p, whereas f_i is the closest embracing site for p.

Now we present an algorithm to compute the E-Voronoi diagram of F using the orthogonal good illumination. We know that the well illuminated points are inside the orthogonal convex hull of F so we start by computing it, uniting at most four monotone chains (see Fig. 3(b)). Afterwards, we decompose the orthogonal convex hull of F by extending horizontal and vertical lines from each light source into the polygon (see Fig. 3(c)). This procedure generates a grid and it can be scanned using the sweeping technique. The resulting partition has a linear number of rays whose arrangement can make up to a quadratic number of rectangles. The algorithm is based on the next lemma.

Lemma 1. *Given a set F of n light sources and a grid that decomposes the orthogonal convex hull of F in rectangles as explained above, every point interior to the same rectangle of the grid shares the light sources' distribution into quadrants.*

Proof. Let F be a set of n light sources and a grid that decomposes the orthogonal convex hull of F into a quadratic number of rectangles as in Fig. 3(c). Suppose that there is an interior point x of a rectangle R which has the light source $f_i \in F$ in some quadrant while another interior point $y \in R$ has f_i in another quadrant. Since the grid is constructed by extending horizontal and vertical lines from each light source into the polygon, one of these lines from f_i must separate x and y into different rectangles. Therefore x and y cannot be interior points to the same rectangle. $\square$

According to this lemma, every point interior to the same rectangle of the grid has the same light sources in the quadrant NE, the same light sources in the quadrant NW, etc. (see Fig. 4(a)). In this subsection, we assume that the points on the border of the rectangles have the same light sources' distribution into quadrants as the interior points. However, this is only true for points of the border of the rectangles that are not simultaneously points of the border

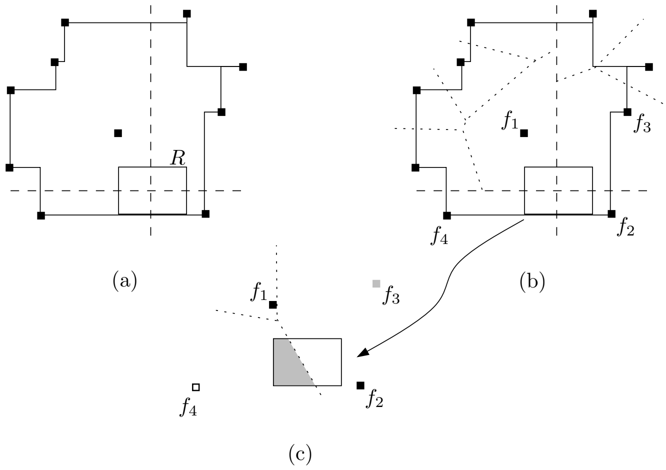

Fig. 4. (a) All the points in R share the light sources' distribution into quadrants. (b) The Voronoi Diagram for the light sources in each quadrant is represent by a dotted line. In this case, all the interior points to R have the same minimal embracing set, $\{f_1, f_2, f_3, f_4\}$. (c) The resulting intersection between R and the Furthest Voronoi Diagram of f_1, f_2, f_3 and f_4 decomposes the rectangle in two regions: E-VR(f_3) (grey region) and E-VR(f_4) (white region).

of the orthogonal convex hull of F. The idea of the algorithm is to compute the E-Voronoi Diagram restricted to each rectangle of the grid and unite them to build the E-Voronoi Diagram of F. For each rectangle R of the grid, we have to compute the points that share their closest embracing site. So we are looking for the points in R that are in the same E-Voronoi region. We compute a usual Voronoi Diagram for the light sources in each of the four quadrants. The intersection of these four Voronoi Diagrams with R gives us the points of R that have the same four closest light sources (one in each quadrant), that is, the points that have the same minimal embracing set (see Fig. 4(b)). So now we compute the points of these regions that share their closest embracing site since it changes according to the light sources' perpendicular bisectors. In order to do this last decomposition of R, we have to compute the Furthest Voronoi Diagram of the four light sources of the minimal embracing set and intersect it with the current region of R (see Fig. 4(c)).

We construct the E-Voronoi Diagram of F repeating this procedure for all the rectangles of the grid and uniting them afterwards (see Fig. 5(a) and 5(b)).

Proposition 3. *Given a set F of n light sources, the described algorithm computes the E-Voronoi Diagram of F in $\mathcal{O}(n^4)$ time.*

Proof. Given a set F of n light sources, computing the orthogonal convex hull of F takes $\mathcal{O}(n \log n)$ time (since it is the union of four monotone chains at the most). To decompose the orthogonal convex hull of F in rectangles we need two

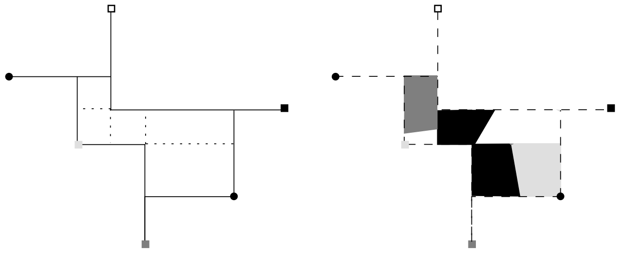

Fig. 5. (a) A set of light sources and its orthogonal convex hull decomposed in rectangles. (b) The E-Voronoi Diagram of the light sources (the light sources represented by a black dot do not have a E-Voronoi region).

sweepings that take $\mathcal{O}(n \log n)$ time though this results in a quadratic number of rectangles. We make a partition of each rectangle in $\mathcal{O}(n^2)$ time by computing its intersection with four Voronoi Diagrams (one per quadrant). For each partition of a rectangle, we intersect it with the Furthest Voronoi Diagram of its minimal embracing set which can be done in $\mathcal{O}(n \log n)$ time. After this procedure, we have computed the E-Voronoi Diagram of F restricted to a rectangle in $\mathcal{O}(n^2)$ time. As we have a quadratic number of rectangles, the union of all these restricted E-Voronoi Diagrams results on the E-Voronoi Diagram of F in $\mathcal{O}(n^4)$ time. □

Once the E-Voronoi Diagram is computed, we can make a query to know exactly where a point is. After the region where the point is has been located, knowing its closest embracing site is straightforward and so is its MER.

4 Good Θ-Illumination

In this section we approach a more general variant of the 1-good illumination of minimum range, the good Θ-illumination. Let F be a set of n light sources in the plane. A cone emanating from a point p is the region between two rays that start at p.

Definition 10. *Let F be a set of n light sources and $\Theta \leq \pi$ a given angle. We say that a point p in the plane is well Θ-illuminated by F if there is, at least, one light source interior to each cone emanating from p with an angle Θ.*

There is an example of this definition in Fig. 6(a) and Fig. 6(b). These well Θ-illuminated points are clearly related to dominance and maximal points. Let $p, q \in S$ be two points in the plane. We say that $p = (p_x, p_y)$ dominates $q = (q_x, q_y)$, $q \prec p$, if $p_x > q_x$ and $p_y > q_y$. Therefore, a point is said to be maximal (or maximum) if it is not dominated or in other words, it means that the quadrant NE centered at p must be empty (see Fig. 6(c)). This version of maximal points can be extended. According to the definition of Avis et. al

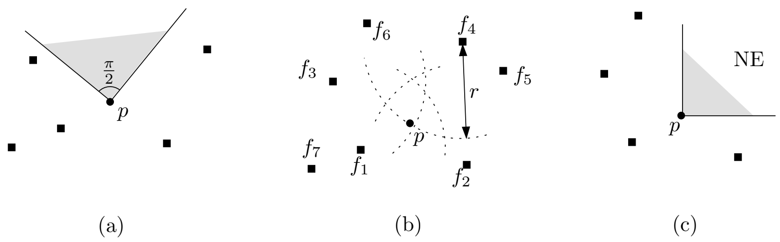

Fig. 6. (a) Point p is not well $\frac{\pi}{2}$-illuminated because there is, at least, one empty cone starting at p with an angle $\frac{\pi}{2}$. (b) Point p is well π-illuminated, its minimal embracing set is $\{f_1, f_2, f_3, f_4\}$ and its MER is r. (c) Point p is a maximum.

[5], a point p in the plane is said to be an unoriented Θ-maximum if there is an empty cone centered at p with an angle of, at least, Θ. The problem of finding all the maximal points of a set S is known as the *maxima problem* [13] and the problem of finding all the unoriented Θ-maximal points is known as the *unoriented* Θ*-maxima problem* [5]. The next proposition follows from the definitions of good Θ-illumination and unoriented Θ-maxima.

Proposition 4. *Let F be a set of n light sources and $\Theta \leq \pi$ a given angle. Given a point p in the plane, p is well Θ-illuminated by F if and only if it is not an unoriented Θ-maximum of the set $F \cup \{p\}$.*

4.1 Minimum Embracing Range of a Well Θ-illuminated Point

We now present a linear time algorithm that not only decides if a point is well Θ-illuminated as it also computes the MER and a minimal embracing set for a given point p in the plane. The main idea of the algorithm is to decide whether a point is well Θ-illuminated by a set of light sources while doing a logarithmic search for its closest embracing site. The logarithmic search is used in the same way as in the algorithm in subsection 2.1, so we will only explain how to decide if a point is well Θ-illuminated by a set of light sources.

Let F be a set of n light sources, p a point in the plane and $\Theta \leq \pi$ a given fixed angle. To check if p is well Θ-illuminated, we divide the plane in several cones of angle $\frac{\Theta}{2}$ emanating from p. Let n_c be the number of possible cones, if 2π is divisible by Θ then $n_c = \frac{4\pi}{\Theta}$ (see Fig. 7(a)). Otherwise $n_c = \lceil \frac{4\pi}{\Theta} \rceil$ because the last cone has an angle less than $\frac{\Theta}{2}$ (see Fig. 7(b)). Since the angle Θ is considered to be a fixed value, the number of cones is constant. Let i be an integer index of arithmetic mod n_c. For $i = 0, \ldots, n_c$, each ray i is defined by the set $\{p + (\cos(\frac{i\Theta}{2}), \sin(\frac{i\Theta}{2}))\lambda : \lambda > 0\}$, while each cone is defined by p and two consecutive rays.

Since we have cones with an angle of at least $\frac{\Theta}{2}$, p is not well Θ-illuminated if we have two consecutive empty cones of angle (see Fig. 7(c)). Note that we have to be sure that the angle of both cones is $\frac{\Theta}{2}$, otherwise this may not be true and we need to proceed as in the third case. If all cones have at least one

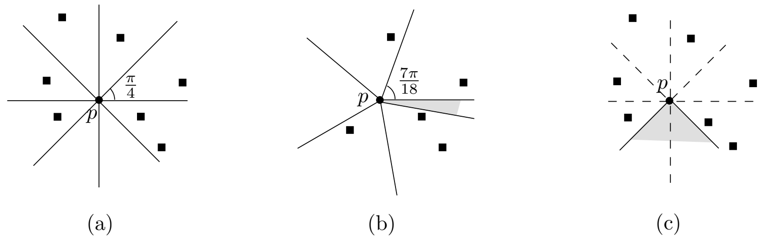

Fig. 7. (a) To check if p is well $\frac{\pi}{2}$-illuminated, the plane is divided in eight cones of angle $\frac{\pi}{4}$. (b) To check if p is well $\frac{7}{9}\pi$-illuminated, the plane is divided in six cones and the last one has an angle less than $\frac{7}{18}\pi$ because 2π is not divisible by $\frac{7}{18}\pi$. (c) Point p is not well $\frac{\pi}{2}$-illuminated because there is an empty cone of angle $\frac{\pi}{2}$.

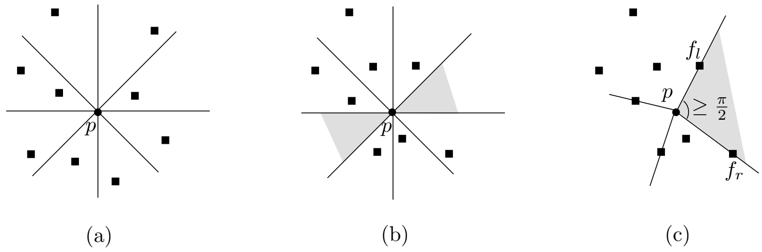

Fig. 8. (a) Point p is well $\frac{\pi}{2}$-illuminated since there is a light source interior to each cone of angle $\frac{\pi}{2}$. (b) There are two non-consecutive empty cones. (c) Point p is not well $\frac{\pi}{2}$-illuminated since there is an empty cone defined by p and the light sources f_l and f_r with an angle greater than $\frac{\pi}{2}$.

interior light source then p is well Θ-illuminated (see Fig. 8(a)). In the last case, there can be at least one empty cone but no two consecutive empty ones (see Fig. 8(b)). We need to spread each empty cone, opening out the rays that define it until we find one light source on each side. Let f_l be the first light source we find on the left and f_r the first light source we find on the right (see Fig. 8(c)). If the angle formed by f_l, p and f_r is at least equal to Θ then there is an empty cone of angle Θ emanating from p. So p is not well Θ-illuminated.

Once the decision algorithm is known, we use it to compute the closest embracing site for p using a logarithmic search. The MER of p is naturally given by the distance between p and its closest embracing site. All the light sources closer to p than its closest embracing site together with the closest embracing site form the minimal embracing set for p. Otherwise p cannot be well Θ-illuminated.

Theorem 1. *Given a set F of n light sources, a point p in the plane and an angle $\Theta \leq \pi$, checking if p is well Θ-illuminated, computing its* MER *and a minimal embracing set for it takes $\Theta(n)$ time.*

Proof. Let F be a set of n light sources, p a point in the plane and $\Theta \leq \pi$ a given angle. Dividing the plane in cones of angle $\frac{\Theta}{2}$ and assigning each light source to its cone takes $\mathcal{O}(n)$ time.

The distances from p to all the light sources can be computed in linear time. Computing the median also takes linear time [6], as well as splitting F in two halves. Since we consider the angle Θ to be a fixed value, the number of cones is constant ($\frac{1}{\Theta}$ is constant). Consequently, spreading each empty cone by computing a light source on each side of the cone is linear. So checking if p is well Θ-illuminated by a set $F' \subseteq F$ is linear on the number of light sources of F'. Note that we never study the same light source twice while searching for the MER of p. So the total time for this logarithmic search is $\mathcal{O}(n + \frac{n}{2} + \frac{n}{4} + \frac{n}{8} + ...) = \mathcal{O}(n)$. Therefore, we compute a closest embracing site and a minimal embracing set for p in linear time.

All the light sources of F are candidates to be the closest embracing site for a point in the plane, so in the worst case we have to study all of them. Knowing this, we have $\Omega(n)$ as a lower bound which makes the linear complexity of this algorithm optimal. □

Note that this algorithm not only computes the minimal embracing set and the MER of a well Θ-illuminated point as it also computes an embracing set for a t-well illuminated point (Definition 4). The next theorem solves the t-good illumination of minimum range using the Θ-illumination of minimum range.

Proposition 5. *Given a set F of n light sources, a point p in the plane and a given angle $\Theta \leq \pi$, let r be the* MER *to well Θ-illuminate p. Then r also t-well illuminates p for $t = \lfloor \frac{\pi}{\Theta} \rfloor$.*

Proof. Let F be a set of n light sources, p a point in the plane and $\Theta \leq \pi$ a given angle. If p is well Θ-illuminated then we know that there is always one interior light source to every cone emanating from p with an angle Θ. On the other hand, p is t-well illuminated if there are, at least, t interior light sources to every half-plane passing through p. An half plane passing through p can be seen as a cone of angle π emanating from p. So if we know that we have at least one light source in every cone of angle Θ emanating from p then we know that we have at least $\lfloor \frac{\pi}{\Theta} \rfloor$ light sources in every half-plane passing through p. This means that p is $\lfloor \frac{\pi}{\Theta} \rfloor$-well illuminated. So the MER needed to well Θ-illuminate p also $\lfloor \frac{\pi}{\Theta} \rfloor$-well illuminates p. □

Corollary 1. *Let F be a set of n light sources, p a point in the plane and $\Theta \leq \pi$ a given angle. A minimal embracing set that well Θ-illuminates p also t-well illuminates p for $t = \lfloor \frac{\pi}{\Theta} \rfloor$.*

Proof. Let F be a set of n light sources, p a point in the plane and $\Theta \leq \pi$ a given angle. According to the last proposition, the MER to well Θ-illuminate p also t-well illuminates it, $t = \lfloor \frac{\pi}{\Theta} \rfloor$. So a closest embracing site for p when it is well Θ-illuminated is at the same distance or further than a closest embracing site for p when t-well illuminated, $t = \lfloor \frac{\pi}{\Theta} \rfloor$. So the minimal embracing set that well Θ-illuminates p also t-well illuminates it for $t = \lfloor \frac{\pi}{\Theta} \rfloor$. □

Fig. 9. (a) Point p is 2-well illuminated since there are at least two light sources in every open half plane passing through p. (b) Point p is not well $\frac{\pi}{2}$-illuminated because there is an empty cone of angle $\frac{\pi}{2}$.

Note 1. If a point is well Θ-illuminated by a set F of light sources, it is also t-well illuminated by F for $t = \lfloor \frac{\pi}{\Theta} \rfloor$, however the other implication is not necessarily true as it is shown in Fig. 9.

5 Conclusions

The visibility problems solved in this paper consider a set of n light sources. Regarding the 1-good illumination, we presented a linear algorithm to compute a Closest Embracing Triangle for a point in the plane and its Minimum Embracing Range (MER). This algorithm can also be used to decide if a point in the plane is 1-well illuminated.

In the following sections, we presented two generalizations of the t-good illumination of minimum range: orthogonal good illumination and the good Θ-illumination of minimum range. We proposed an optimal linear time algorithm to compute the MER of an orthogonally well illuminated point, as well as its minimal embracing set. Related to this variant, the E-Voronoi Diagram was also presented as well as an algorithm to compute it that runs in $\mathcal{O}(n^4)$ time.

We introduced the Θ-illumination of minimum range and an optimal linear time algorithm. The algorithm computes the MER needed to well Θ-illuminate a point in the plane and a minimal embracing set for it. We established a connection between the t-good illumination of minimum range and the good Θ-illumination of minimum range in Proposition 5. The MER to well Θ-illuminate a point also t-well illuminates that point, for $t = \lfloor \frac{\pi}{\Theta} \rfloor$.

All the algorithms in this paper apart from the one that computes the E-Voronoi Diagram have been implemented using Java. They all have been implemented without any major issues and take the expected run-time to compute a solution. The algorithm to compute the E-Voronoi Diagram is by far the most challenging since it needs a good data structure to compute and merge five Voronoi Diagrams. Nevertheless, this must be done a quadratic number of times which can be disastrous if the data structure takes too much time to be processed. Though it hasn't been implemented yet, it is in our plans to do so.

Acknowledgments We wish to thank Belén Palop from the Universidad de Valladolid for helpful discussions about the algorithm to compute the MER of a 1-well illuminated point.

References

1. Abellanas, M., Bajuelos, A., Hernández, G., Matos, I.: Good Illumination with Limited Visibility. Proceedings of the International Conference of Numerical Analysis and Applied Mathematics, Wiley-VCH Verlag (2005) 35–38.
2. Abellanas, M., Bajuelos, A., Hernández, G., Matos, I., Palop, B.: Minimum Illumination Range Voronoi Diagrams. Proceedings of the 2^{nd} International Symposium on Voronoi Diagrams in Science and Engineering (2005) 231–238.
3. Abellanas, M., Canales, S., Hernández, G.: Buena iluminación. Actas de las IV Jornadas de Matemática Discreta y Algorítmica (2004) 239–246.
4. Asano, T., Ghosh, S. K., Shermer, T. C.: Visibility in the plane. In: Sack, J.-R., Urrutia, J. (eds): Handbook of Computational Geometry. Elsevier (2000) 829–876.
5. Avis, D., Beresford-Smith, B., Devroye, L., Elgindy, H., Guévremont, H., Hurtado, F., Zhu, B.: Unoriented Θ-maxima in the plane: complexity and algorithms. Siam J. Computation, Vol. 28, No. 1, (1998) 278–296.
6. Blum, M., Floyd, R.W., Pratt, V., Rivest, R., Tarjan, R.: Time bounds for selection. Journal of Computer and System Sciences **7** (1973) 448–461.
7. Canales, S.: Métodos heurísticos en problemas geométricos, Visibilidad, iluminación y vigilancia. Ph.D. thesis, Universidad Politécnica de Madrid (2004).
8. Chan, M.Y., Chen, D., Chin, F. Y. L., Wang, C. A.: Construction of the Nearest Neighbor Embracing Graph of a Point Set. Journal of Combinatorial Optimization, Vol. 11, No. 4, (2006), 435–443.
9. Chiu, S. N., Molchanov, I. S.: A new graph related to the directions of nearest neighbours in a point process. Advances in Applied Probability, Vol. 35, No. 1, (2003) 47–55.
10. Efrat, A., Har-Peled, S., Mitchell, J. S. B.: Approximation Algorithms for Two Optimal Location Problems in Sensor Networks. Proceedings of the 14^{th} Annual Fall Workshop on Computational Geometry, MIT (2004).
11. Karlsson R., Overmars, M.: Scanline Algorithms on a Grid. BIT Numerical Mathematics, Vol. 28, No. 2, Springer Netherlands (1988) 227–241.
12. Megiddo, N.: Linear-time algorithms for linear programming in $\mathbb{R}^3$ and related problems. SIAM Journal on Computing, Vol. 12, No. 4, (1983) 759–776.
13. Kung, H., Luccio, F., Preparata, F.: On finding the maxima of a set of vectors. Journal of ACM **22** (1975) 469–476.
14. Ntafos, S.: Watchman routes under limited visibility. Computational Geometry: Theory and Applications, Vol. 1, No. 3, Elsevier (1992) 149–170.
15. Smith, J., Evans, W.: Triangle Guarding. Proceedings of the 15^{th} Canadian Conference on Computational Geometry (2003) 76–80.
16. Urrutia, J.: Art Gallery and Illumination Problems. In: Sack, J.-R., Urrutia, J. (eds): Handbook of Computational Geometry. Elsevier (2000) 973–1027.