Variations of Good Illumination*

Manuel Abellanas[†] Antonio Bajuelos[‡] Inês Matos[‡]

Abstract

A point p is 1-well illuminated by a set F of n point lights if there is, at least, one light source interior to each half plane with p on its border (also known as $\triangle$-guarding or well-covering). We present two generalizations of 1-good illumination: the orthogonal good illumination and the good Θ-illumination. For the first, we propose a linear time algorithm to optimize the light sources' illumination range to orthogonally well illuminate a point, as well as minimal embracing set for it. We follow presenting the E-Voronoi Diagram for this variant and an algorithm to compute it that runs in $\mathcal{O}(n^4)$ time. A point p in the plane is well Θ-illuminated if there is, at least, one light source interior to each cone emanating from p of angle Θ. Given a fixed angle $\Theta \leq \pi$, we present an algorithm to minimize the light sources' illumination range to well Θ-illuminate a point p. The algorithm runs in linear time which is optimal and it also computes a minimal embracing set of light sources for p. Good Θ-illumination is related to the concept of t-good illumination as it is shown in Proposition 3.4.

1 Introduction and Related Works

Visibility and illumination have been a main topic for different papers in the area of Computational Geometry. However, most of these problems cannot be applied to real life, since they deal with ideal concepts. For instance, light sources have some restrictions since they cannot illuminate an infinite region as their light naturally fades as the distance grows. As well as cameras or robot vision systems, both have severe visibility range restrictions because 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 a limited illumination range so their illuminated regions are delimited. We use a limited visibility definition due to Ntafos [11] as well as a concept related to this type of problems, the t-good illumination due to Canales et. al [3, 6]. This paper is focused in an optimization problem related to limited range illumination.

This paper is structured as follows. In the next subsection we formalize the 1-good illumination and define the problem of optimizing the light sources' illumination range. Sections 2 and 3 are devoted to extensions of 1-good illumination. In section 2 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 3 we extend 1-good illumination to cones and make a brief relation between this variant and the Maxima Problem [4, 10]. We conclude this paper in section 4.

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 [7]. The set CH(F) represents the convex hull of the set F.

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

[†]Facultad de Informática, Universidad Politécnica de Madrid, supported by grant TIC2003-08933-C02-01, MEL-HP2005-0137 and partially supported by CAM:P-DPI-000235-0505 `mabellanas@fi.upm.es`

[‡]Departamento de Matemática & CEOC, Universidade de Aveiro, supported by CEOC through *Programa* POCTI, FCT, co-financed by EC fund FEDER and by Acção No. E-77/06 `{leslie,ipmatos}@mat.ua.pt`

Definition 1.1. 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 CH(F).

Definition 1.2. A site $f_i \in F$ is an embracing site for 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 Figure 1(a) and Figure 1(b)).

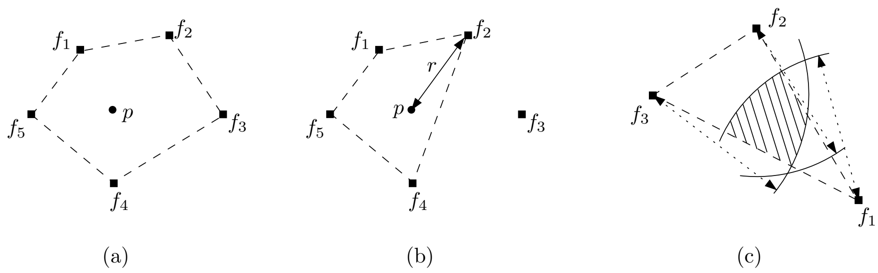

Figure 1: (a) The light sources f_2 and f_3 are embracing sites for point p. (b) The light source f_2 is the closest embracing site for p and its illumination range is $r = d(p, f_2)$. The set $\{f_1, f_2, f_4, f_5\}$ is a minimal embracing set for p. (c) $A_r^E(f_1, f_2, f_3)$ is the shaded open area,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

Definition 1.3. 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 1.4 ([6])**.** 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. The 1-good illumination can also be found under the name of $\triangle$-guarding [12] or well-covering [8]. 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 [8]. 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 $A_r^E(f_i, f_j, f_k) = A_r(f_i, f_j, f_k) \cap \mathrm{int}(\mathrm{CH}(f_i, f_j, f_k))$ to denote the illuminated area embraced by the light sources f_i, f_j and f_k.

Definition 1.5. 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 1.6. 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 MER(F, p) or MER(F, S).

Figure 1(c) illustrates definition 1.5. Since the set F is clear from the context, we will use "MER of p" instead of MER(F, p) and "MER of S" instead of 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.

2 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 Figure 2(a).

Definition 2.1. 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, 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 is the orthogonal convex hull (see Karlsson and M. Overmars [9]). 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. If we connect each pair of interior points using only horizontal and vertical line segments, the polygonal line between them lies inside the polygon. 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 Figure 2(b)).

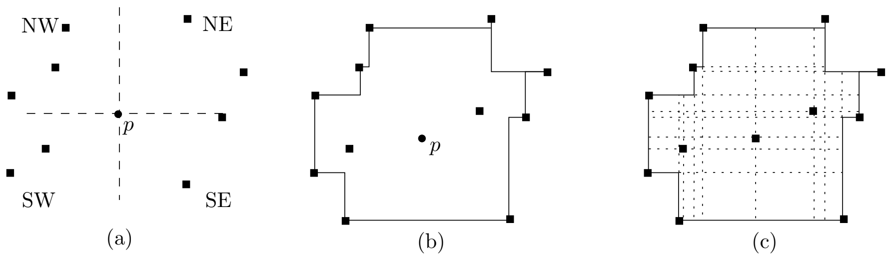

Figure 2: (a) Point p is orthogonally well illuminated because all oriented quadrants centered at p are non-empty. (b) Point p is interior to the orthogonal convex hull of F, so it is orthogonally well illuminated. (c) An orthogonal convex hull decomposed into several rectangles.

2.1 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 (the closest light source to p in each quadrant). From these four light sources, the closest embracing site for p is the furthest. The MER of p is given by the distance between p and its closest embracing site.

Proposition 2.2. *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. □

2.2 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 [2] when considering the usual 1-good illumination.

Definition 2.3 ([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 $\text{E-VR}(f_i, F) = \{x \in \mathbb{R}^2 : f_i \text{ is the closest embracing site for } x\}$.

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 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 Figure 2(b)). Afterwards, we decompose the orthogonal convex hull of F by extending orthogonal rays from each light source into the polygon (see Figure 2(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 2.4. *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 Figure 2(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 orthogonal rays from each light source into the polygon, one of the orthogonal rays from f_i must separate x and y into different rectangles. Therefore x and y cannot be interior points to the same rectangle. ☐

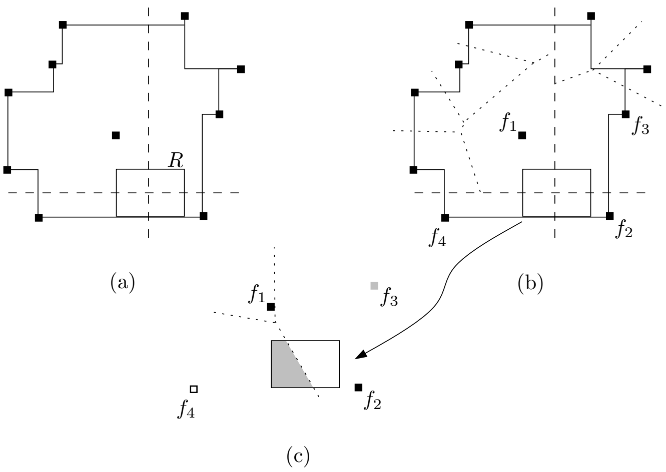

Figure 3: (a) All the points in R share the light sources' distribution into quadrants. (b) The Voronoi Diagram for the light sources in each quadrant is represent by a dotted line. In this case, all the interior points to R have the same minimal embracing set, $\{f_1, f_2, f_3, f_4\}$. (c) The resulting intersection between R and the Furthest Voronoi Diagram of f_1, f_2, f_3 and f_4 decomposes the rectangle in two regions: $\text{E-VR}(f_3)$ (grey region) and $\text{E-VR}(f_4)$ (white region).

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 Figure 3(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 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. Next we intersect R with four Voronoi Diagrams (one for each quadrant) and decompose it in several regions. The points in each region have the same minimal embracing set (see Figure 3(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 Figure 3(c)). We construct the E-Voronoi Diagram of F repeating this procedure for all the rectangles of the grid and uniting them afterwards

Proposition 2.5. *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 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. $\square$

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.

3 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 3.1. 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.

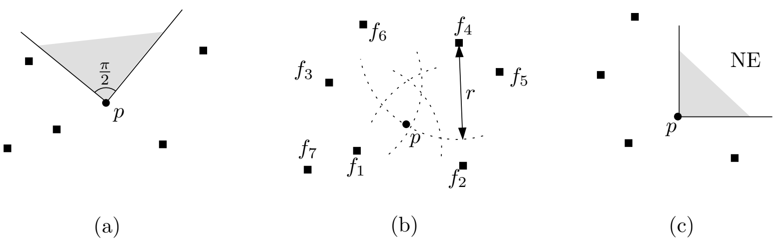

Figure 4: (a) Point p is not well $\frac{\pi}{2}$-illuminated because there is, at least, one empty cone starting at p with an angle $\frac{\pi}{2}$. (b) Point p is well π-illuminated, its minimal embracing set is $\{f_1, f_2, f_3, f_4\}$ and its MER is r. (c) Point p is a maximum.

There is an example of this definition in Figure 4(a) and Figure 4(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 Figure 4(c)). This version of maximal points can be extended. According to the definition of Avis et. al [4], 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* [10] and the problem of finding all the unoriented Θ-maximal points is known as the *unoriented Θ-maxima problem* [4]. The next proposition follows from the definitions of good Θ-illumination and unoriented Θ-maxima.

Proposition 3.2. *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 $F \cup \{p\}$.*

3.1 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.

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 Figure 5(a)). Otherwise $n_c = \lceil \frac{4\pi}{\Theta} \rceil$ because the last cone has an angle less than $\frac{\Theta}{2}$ (see Figure 5(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.

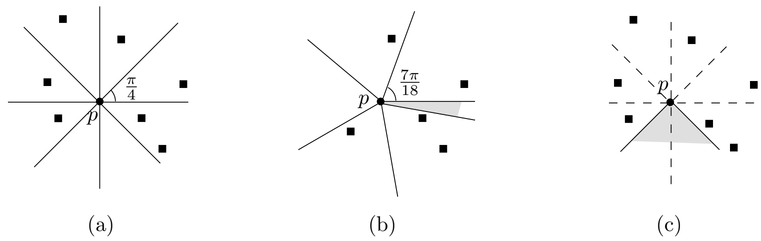

Figure 5: (a) To check if p is well $\frac{\pi}{2}$-illuminated, the plane is divided in eight cones of angle $\frac{\pi}{4}$. (b) To check if p is well $\frac{7}{9}\pi$-illuminated, the plane is divided in six cones and the last one has an angle less than $\frac{7}{18}\pi$ because 2π is not divisible by $\frac{7}{18}\pi$. (c) Point p is not well $\frac{\pi}{2}$-illuminated because there is an empty cone of angle $\frac{\pi}{2}$.

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 Figure 5(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 interior light source then p is well Θ-illuminated (see Figure 6(a)). In the last case, there can be at least one empty cone but no two consecutive empty ones (see Figure 6(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 Figure 6(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 a logarithmic search to compute the MER of p. First, we compute the median of the distances between the light sources and p and divide F in two subsets. The subset F_c contains the $\lceil \frac{n}{2} \rceil$ closest light sources to p, while the set F_f contains the furthest half. Using the method described above, we are able to decide if the light sources of F_c are enough to well Θ-illuminate p. If they are then we can forget about the light sources of F_f and compute the median of the distances between the light sources of F_c and p. We reassign F_c to the closest half and repeat the previous method to check if the new set is still enough to well Θ-illuminate p. If F_c does not well Θ-illuminate p then we have to save the location of the empty cones. Since the light sources of F_c are not enough, we have to get some more from the set F_f that have not been used. We compute the median of the light sources of F_f and reassign F_c to the closest half of the latter. Now we will test

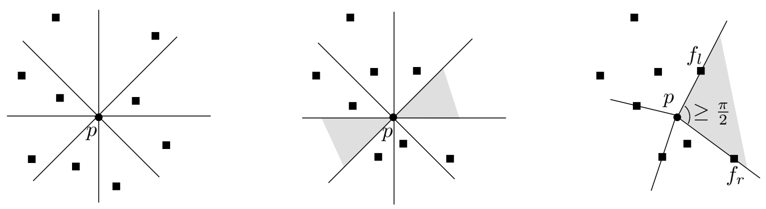

(a) (b) (c)

Figure 6: (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 $\frac{\pi}{2}$.

the new set F_c just by checking if its new light sources are interior to the cones left empty in the last iteration. If they are not enough, then we repeat the procedure and try again with another half of F_f, otherwise we try half from the ones we have just used. We repeat this search until $|F_c| = 1$. If p is well Θ-illuminated then the only light source of F_c is the closest embracing site for p. 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 3.3. *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 [5], 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. $\square$

Note that this algorithm not only computes the minimal embracing set and the MER of a well Θ-illuminated point as it also computes them for a t-well illuminated point (Definition 1.4). The next proposition solves the t-good illumination of minimum range using the Θ-illumination of minimum range.

Proposition 3.4. *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. $\square$

Corollary 3.5. *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 $t = \lfloor \frac{\pi}{\Theta} \rfloor$.*

Remark 3.6. 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.

4 Conclusions

The visibility problems solved in this paper consider a set of n light sources. 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 3.4. The MER to well Θ-illuminate a point also t-well illuminates that point, for $t = \lfloor \frac{\pi}{\Theta} \rfloor$.

References

[1] M. Abellanas, A. Bajuelos, G. Hernández, I. Matos, Good Illumination with Limited Visibility. *Proceedings of the International Conference of Numerical Analysis and Applied Mathematics*, Wiley-VCH Verlag, 35–38, 2005.

[2] M. Abellanas, A. Bajuelos, G. Hernández, I. Matos and B. Palop, Minimum Illumination Range Voronoi Diagrams. *Proceedings of the 2^{nd} International Symposium on Voronoi Diagrams in Science and Engineering*, 231–238, 2005.

[3] M. Abellanas, S. Canales and G. Hernández: Buena iluminación. *Actas de las IV Jornadas de Matemática Discreta y Algorítmica*, 239–246, 2004.

[4] D. Avis, B. Beresford-Smith, L. Devroye, H. Elgindy, H. Guévremont, F. Hurtado, and B. Zhu, Unoriented Θ-maxima in the plane: complexity and algorithms. *Siam J. Computation* **28**(1), 278–296, 1998.

[5] M. Blum, R. W. Floyd, V. Pratt, R. Rivest and R. Tarjan, Time bounds for selection. *Journal of Computer and System Sciences* **7**, 448–461, 1973.

[6] S. Canales, Métodos heurísticos en problemas geométricos, visibilidad, iluminación y vigilancia. Ph.D. thesis, Universidad Politécnica de Madrid (2004).

[7] S. N. Chiu and I. S. Molchanov, A new graph related to the directions of nearest neighbours in a point process. *Advances in Applied Probability* **35**(1), 47–55, 2003.

[8] A. Efrat, S. Har-Peled and J. S. B. Mitchell, Approximation Algorithms for Two Optimal Location Problems in Sensor Networks. *Proceedings of the 14^{th} Annual Fall Workshop on Computational Geometry*, MIT, 2004.

[9] R. Karlsson and M. Overmars, Scanline Algorithms on a Grid. *BIT Numerical Mathematics* **28**(2), Springer Netherlands, 227–241, 1988.

[10] H. Kung, F. Luccio, and F. Preparata, On finding the maxima of a set of vectors. *Journal of ACM* **22**, 469–476, 1975.

[11] S. Ntafos, Watchman routes under limited visibility. *Computational Geometry: Theory and Applications* **1**(3), Elsevier, 149–170, 1992.

[12] J. Smith and W. Evans, Triangle Guarding. *Proceedings of the 15^{th} Canadian Conference on Computational Geometry*, 76–80, 2003.