Available online at www.sciencedirect.com
ScienceDirect
European Journal of Operational Research 181 (2007) 20–29
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/ejor

Continuous Optimization

The maximin line problem with regional demand

J.M. Díaz-Báñez [a,*,1], P.A. Ramos [b,2], P. Sabariego [c,3]

[a] Departamento de Matemática Aplicada II, Universidad de Sevilla, Escuela Superior de Ingenieros, Camino de los Descubrimientos s/n, 41092 Sevilla, Spain
[b] Escuela Politécnica, E-226, Universidad de Alcalá, Aptdo, de Correos 20, 28871 Alcalá de Henarts, Madrid, Spain
[c] Dpto de Matematicas, Estadistica y Computación, Facultad de Ciencios, Avda de los Castros s/n, 39005 Cantabria, Spain

Received 23 September 2005; accepted 7 June 2006
Available online 14 August 2006

Abstract

Given a family $\mathscr{P} = \{P_1, \ldots, P_m\}$ of m polygonal regions (possibly intersecting) in the plane, we consider the problem of locating a straight line ℓ intersecting the convex hull of $\mathscr{P}$ and such that $\min_k d(P_k, \ell)$ is maximal. We give an algorithm that solves the problem in time $O((m^2 + n\log m)\log n)$ using $O(m^2 + n)$ space, where n is the total number of vertices of $P_1, \ldots, P_m$. The previous best running time for this problem was $O(n^2)$ [J. Janardan, F.P. Preparata, Widest-corridor problems, Nordic Journal of Computing 1 (1994) 231–245]. We also improve the known complexity for several variants of this problem which include a three dimensional version – the maximin plane problem –, the weighted problem and considering measuring distance different to the Euclidean one.

Keywords: Location; Computational geometry; Linear facility; Duality

1. Introduction

The advances in computational geometry have given rise to the development of efficient algorithms to solve facility location problems. In fact, since its beginning, there has been a strong interaction between both fields. A proof of this is the fact that one of the problems which were solved in the origins of computational geometry, computing the *minimum spanning circle*, is a geometric interpretation of the best known facility location problem, the *one-center problem*.

This interaction has resulted in a wealth of papers and results of interest to researchers and practitioners in both fields. In this sense, the surveys [29,10,22] establish the current state-of-the-art for single and non-single facility location problems, respectively.

Linear facility location has been of great interest both in location theory [25,26,33,3] and in computational geometry [23,20,22]. In this paper we deal with the placement of an undesirable facility modelled by a line, amidst polygonal regions. The computation of obnoxious routes has become a topic of

* Corresponding author.
E-mail addresses: dbanez@us.es (J.M. Díaz-Báñez), pedro.ramos@uah.es (P.A. Ramos), pilar.sabariego@unican.es (P. Sabarieg.
1 Partially supported by Project BFM2003-04062.
2 Partially supported by Projects TIC2003-08933-C02-01 and BMF2002-04402-C02-01.
3 Supported in part by Project BFM2001-1153.

0377-2217/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2006.06.013

increasing interest in recent years. In fact, there is a natural reason to obtain maximum 'clearance' in many applications as the design of channels for transportation of hazardous materials or paths avoiding obstacles in robotics [7,8].

The problem of locating a linear route which maximizes the minimum weighted Euclidean distance to a set of points was first considered in [12]. In this paper, a naive $O(n^3)$-time algorithm was proposed. However, by using topological sweeping and duality, the unweighted version of this problem can be solved in $O(n^2)$ time and $O(n)$ space [19]. In fact, Houle and Maciel address in [19] the problem of computing a widest empty corridor through a set of points in the plane, which is precisely an equivalent formulation of the maximin line problem.

On the other hand, although in the classical facility location problems the existing facilities are represented as a set of points, there exists a real interest in considering models involving regions as demand sites. In this way, different real world situations can be modelled better than the classical versions [11,27]. In this case, the distance between the facility and the customer may be calculated as some form of expected or average travel distance, for instance, see [5], or the distance to the closest point on the boundary of the region [4].

The problem of locating an obnoxious line in presence of polygonal regions was considered in [18,17]. A brute-force $O(n^4)$-time algorithm was proposed both with Euclidean and polyhedral norms. However, because the problem is just the empty corridor problem, an efficient $O(n^2)$-time algorithm can be found in [21] for the unweighted Euclidean case. Given a set $\mathscr{P} = \{P_1, \ldots, P_m\}$ of m polygonal regions in the plane with a total of n edges, we wish to compute efficiently a maximin line with respect to $\mathscr{P}$ for a general case. Since in typical applications $m \ll n$, we would like to have an algorithm whose performance depends both on m and n and which is significantly better than the previous one when $m \ll n$. Our results include:

1. An algorithm to compute the unweighted maximin line through $\mathscr{P}$, both for the Euclidean metric and for a general metric, in $O((m^2 + n\log m)\log n)$ time and $O(m^2 + n)$ space. This result improves the $O(n^2)$-time algorithm of [21] if $m \ll n$.
2. The adaptation of the method to solve the general case in which the polygonal regions are weighted (the weights represent the number of inhabitants, for instance). This variant of the problem is solved in $O(nm + n\log^2 n\log m + m^2\log n\log^2 m)$ time and $O(m^2 + n)$ space, which significantly improves the bound $O(n^4)$ of [18,17].
3. An algorithm to solve an extension of the Euclidean unweighted version of the problem to three dimensions, the *obnoxious plane problem in presence of polyhedra*, in $O(m^2 n\log(m^2 n))$ time and $O(m^2 n)$ space. This bound improves on the $O(n^3)$ time proposed in [9].

The rest of the paper is organized as follows. In Section 2 we state the problem, present some geometric preliminaries and briefly describe known computational results. Our general approach is proposed in Section 3 for the Euclidean unweighted case. In Section 4 the method is adapted to solve a more general model. Section 5 address the three-dimensional extension of the problem.

2. Overview

We start by introducing some notation and considering a few geometric preliminaries. A summary of related results is also presented. Let $\mathscr{P} = \{P_1, \ldots, P_m\}$ be a set of m polygonal regions in the plane with a total of n vertices. The distance between a polygon $\mathscr{P}$ and a straight line ℓ is given by the shortest distance based on the distance measure d, i.e., $d(P, \ell) = \min_{p\in P, x\in\ell} d(p, x)$, where $d(p, x)$ denotes the Euclidean distance between points p and x. First, we address the Euclidean case and then we adapt the approach to solve other versions, including weighted and arbitrary distances.

In the definition of an obnoxious facility location problem, the location of the facility must be constrained, as otherwise it may be simply removed to infinity. The facility is normally constrained to go through some sort of bounding region. As in [18], we are looking for a linear route inside the convex hull of the set $\mathscr{P}$.

The *maximin line through* $\mathscr{P}$ *problem* can be formalized as follows:

Given a set $\mathscr{P}$ of m (possibly intersecting and non-convex) polygons with a total of n vertices, compute a line ℓ such that

1. $\min_{P\in\mathscr{P}} d(P, \ell)$ is maximal, and
2. ℓ divides $\mathscr{P}$ into two non-empty subsets.

Because the closest point of a polygon P to a line ℓ not intersecting P is always a vertex of the convex hull of P, hereafter we consider that polygons are convex and, if this is not the case, we compute their convex hull as a preliminary step.

The problem can be reformulated equivalently as the computation of the widest empty corridor through polygonal obstacles [21]. In this geometric formulation, an empty corridor C, through $\mathscr{P}$, is the open region of the plane that is enclosed by two parallel straight lines intersecting the convex hull of $\mathscr{P}$ and such that the region does not intersect any polygon in $\mathscr{P}$.

Let us observe that a given set $\mathscr{P}$ may have no empty corridor through it and, if this is the case, the maximin line problem has no solution. Therefore, *decision* (deciding whether or not there exists a line ℓ through $\mathscr{P}$ not intersecting any polygon) and *optimization* (finding the farthest one) problems can be independently considered. This fact suggests trying to compute first a feasible set of lines and then find the optimal solution in that set.

The following lemma characterizes the solution to our problem. The proof is the same as in [19], where an analogous result is presented for the case of points instead of polygons.

Lemma 1. *Let ℓ^* be an optimal line and let ℓ_1 and ℓ_2 be the bounding lines of the corridor generated by ℓ^*. Then, one of the following conditions must hold*:

(a) *ℓ_1 and ℓ_2 contain vertices v_1 and v_2, and the lines are perpendicular to the line segment connecting v_1 and v_2.*
(b) *There are two vertices on ℓ_1 and one vertex on ℓ_2 (or the opposite) and, furthermore, the vertex on ℓ_1 is between the vertices on ℓ_1 when viewed from a direction orthogonal to ℓ^*.*

This lemma guarantees an $O(n^3)$ upper bound on the number of candidate lines for the two types of corridors. In [18,17], a straightforward $O(n^4)$-time algorithm is proposed by exhaustively considering all possible cases and finding the optimal one. A more efficient algorithm was proposed in [21]. For the sake of completeness, we include here the approach as well.

The main idea is to interpret conditions (a) and (b) of Lemma 1 in the dual plane by using the duality transformation mapping the non-vertical line ℓ with equation $y = mx - n$ to its dual point $\ell^* = (m,n)$ and the point $p = (a,b)$ to its dual line p^*: $y = ax - b$. By using dual transformation properties, it follows that if pq is a segment, the dual of the set of lines intersecting pq is the double wedge defined by lines p^* and q^* and not containing the vertical line.

Let $\mathscr{H}$ be the set of lines dual to vertices of $\mathscr{P}$ and let $\mathscr{A}(\mathscr{H})$ be the arrangement in the plane induced by $\mathscr{H}$. The properties of the duality transform can be used to characterize in $\mathscr{A}(\mathscr{H})$ the sets of type-(a) and type-(b) corridors of Lemma 1.

A corridor C with bounding lines ℓ_1 and ℓ_2 is represented in the dual plane by the vertical segment with endpoints ℓ_1^* and ℓ_2^*. If C is an type-(a) empty corridor, then C corresponds to a vertical segment inside a face of $\mathscr{A}(\mathscr{H})$ that connects a vertex and a edge of that cell. Similarly, a type-(b) empty corridor corresponds to a vertical segment inside a cell connecting an edge with an edge. In this case, the uniqueness of the segment follows from the perpendicularity condition.

Furthermore, the set of lines intersecting a given set of edges of $\mathscr{P}$ corresponds to a face of $\mathscr{A}(\mathscr{H})$ and, in particular, lines avoiding all the polygons of $\mathscr{P}$ correspond to some set of faces of the arrangement. As observed in [21], the topological sweep of [13] can be adapted to compute such set of faces and therefore we have:

Theorem 1 [21]. *An optimal maximin line for the set of polygons $\mathscr{P}$ can be computed in $O(n^2)$ time and $O(n)$ space, where n is the total number of vertices of the polygons.*

3. Our approach

In this section we show that if m is small compared to n it is more convenient to avoid the construction of the whole arrangement of lines dual to vertices of $\mathscr{P}$. We show that the set of lines not intersecting any polygon in $\mathscr{P}$ has complexity $O(m^2 + n)$ and can be constructed in time $O((m^2 + n\log m)\log n)$. Therefore, if $m = o(n)$, the time complexity is improved perhaps with some extra memory cost, while if $m = O(\sqrt{n})$ time complexity is improved with the same space complexity.

Given a polygon $P \in \mathscr{P}$, let us denote by U_P and L_P the dual sets of the lines above P and bellow P, respectively. It is well known that U_P and L_P are disjoint convex polygons (one unbounded from above, the other unbounded from below), as shown in Fig. 1. Then, $\bar{P}^* = U_P \cup L_P$ is the set of points dual to the lines not intersecting P. A vertex of $\bar{P}^*$ is the dual of a line supporting an edge of P.

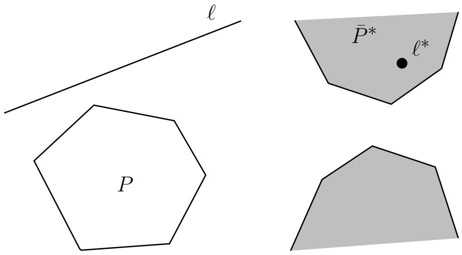

Fig. 1. Dual interpretation of a free-collision line.

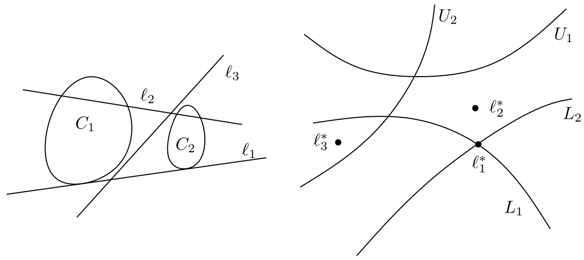

Fig. 2. Duality for convex sets P_1 and P_2.

We are interested in the set $\mathscr{A}^* = \bigcap_{P \in \mathscr{P}} \bar{P}^*$ because a point $\ell^* \in \mathscr{A}^*$ corresponds, in primal space, to a line ℓ that does not intersect any polygon in $\mathscr{P}$.

The complexity of this set is defined as the sum of its vertices, edges and faces and is proportional to the number of vertices. A vertex v of $\mathscr{A}^*$ is either a vertex of some polygonal region $\bar{P}^*$ or a point dual to a common tangent of two polygons in $\mathscr{P}$ that does not intersect any other polygon. Clearly, there are at most n vertices of the first type. In [1] it is shown that the number of vertices of the second type is $O(m^2 + n)$ using an argument similar to the following one. Let $Q_1, Q_2, \ldots, Q_t$, where $t \leqslant m$, be the connected components of $\bigcup_{P \in \mathscr{P}} P$. A vertex of the second type corresponds either to a line tangent to some Q_i – there are $O(n)$ of those – or to a common tangent of Q_i and Q_j. Because two disjoint polygons have at most four common tangents, we conclude that the number of vertices of the second type is $O(m^2 + n)$. In the same paper it is shown that a simple divide-and-conquer algorithm which performs the conquer approach doing a sweep computes the set $\mathscr{A}^*$ in time $O((m^2 + n \log m) \log n)$. We summarize this discussion in the following result.

Lemma 2. *$\mathscr{A}^*$ has complexity $O(m^2 + n)$ and can be computed in time $O((m^2 + n \log m) \log n)$.*

Once $\mathscr{A}^*$ is computed, the decision problem reduces to check whether it has some face which correspond to a line having polygons on both sides, i.e., a face of the arrangement which has edges both above and bellow it. The optimization problem can be solved visiting all the faces of $\mathscr{A}^*$: within each face, we have to compute the width of the slabs which correspond in the dual plane to vertical segments connecting either two edges of $\mathscr{A}^*$ (type(a)-corridor) or a vertex and an edge (type(b)-corridor), as defined in Lemma 1. Clearly, this can be done in time proportional to the size of the face performing a sweep of the face. Therefore, we get the following result:

Theorem 2. *Let $\mathscr{P}$ be a set of m convex polygons with a total of n vertices. An optimal maximin line trough $\mathscr{P}$ can be found in time $O((m^2 + n \log m) \log n)$ using $O(m^2 + n)$ space.*

The approach in this section can be applied to the case when the obstacles are convex sets with constant description complexity, i.e., whose boundaries are algebraic curves with degree bounded by a certain constant. Again, we use duality and observe that the set of lines intersecting a convex set P corresponds in dual plane to the region between two convex curves, U and L, which are dual to the set of upper and lower tangents of P, respectively (see Fig. 2). Therefore, the set of lines avoiding a family of convex sets $P_1, P_2, \ldots, P_m$ corresponds to a set of faces in the arrangement formed by those curves. The crucial parameter which bounds the complexity of the arrangement is the number of intersections between any two curves. If every two curves intersect in at most one point, then the complexity of the arrangement is $O(m^2)$ and can be constructed within the same asymptotic time (see [32]). Once the arrangement is constructed, the problem of computing the widest empty corridor can be easily solved. We observe that the condition for the number of intersections corresponds in primal plane to the fact that every two convex sets P_i and P_j have at most one common upper tangent, one common lower tangent and two inner tangents and is satisfied in a variety of situations, for instance, if the convex sets are pairwise disjoint or if we are dealing with a family of arbitrary disks.

4. The weighted maximin line problem with arbitrary norms

In this section we generalize the problem both by considering distances different from the Euclidean and by adding weights to the sites. Facility location models usually consider given weights associated to the input, representing the importance of the existing facilities. Also, non-Euclidean norms to measure distances have been widely used in the literature [16,33]. Let us start by introducing some notation that we borrow from [33].

Let $\mathscr{P} = \{P_1, \ldots, P_m\}$ be the family of convex polygonal regions and let w_k be the weight associated to the polygon P_k. Let $\mathscr{B}$ be a convex, compact, centrally symmetric set in the plane. The norm with unit ball $\mathscr{B}$ is defined by $\gamma\mathscr{B}(x) := \min\{|\lambda| : x \in \lambda\mathscr{B}\}$. The induced distance between two points x and y is denoted $d_{\mathscr{B}}(x,y) := \gamma_{\mathscr{B}}(x-y)$. If A and B are two closed subsets in $\mathbb{R}^2$, then the distance between A and B is defined as $d_{\mathscr{B}}(A,B) := \min_{a\in A, b\in B} d_{\mathscr{B}}(a,b)$.

We consider both the decision and the optimization versions of the *weighted maximin line problem with arbitrary norm*. If we denote by $\mathscr{L}_P$ the set of lines in the plane intersecting $CH(\mathscr{P})$, the problems can be defined, respectively, as follows:

[Dec] Given $\delta > 0$, does there exist line $\ell \in \mathscr{L}_P$ such that $\min_{P_k\in\mathscr{P}} \frac{1}{\omega_k} d_{\mathscr{B}}(P_k, \ell) \geqslant \delta$?
[Opt] Compute $\max_{\ell\in\mathscr{L}_P} \min_{P_k\in\mathscr{P}} \frac{1}{\omega_k} d_{\mathscr{B}}(P_k, \ell)$.

Our plan is to give an efficient solution to the decision problem and then applying parametric search in order to solve the optimization problem. First we recall the concept of Minkowski sum: given two sets $A, B \subset \mathbb{R}^2$, the Minkowski sum of A and B is defined as

$$A \oplus B = \{a + b | a \in A, b \in B\},$$

where a and b are added up because they are interpreted as vectors once an origin has been fixed. Because

$$\frac{1}{\omega_k} d_{\mathscr{B}}(P_k, \ell) \geqslant \delta \iff \ell \cap (P_k \oplus \omega_k\delta\mathscr{B}) = \varnothing \qquad (1)$$

the decision problem can be formulated in the following way: given $\delta > 0$, decide whether there exists a line $\ell \in \mathscr{L}_P$ that does not intersect the sets $P_i \oplus \omega_1\delta\mathscr{B}$, for $i = 1, \ldots, m$ (see Fig. 3).

At this point we have to make some assumptions on the ball $\mathscr{B}$ that allows us to perform computations. For instance, we can assume that $\mathscr{B}$ is a convex polygon with a constant number of edges[4] or an algebraic curve with constant description complexity. In the later case, we assume that our model of computation is powerful enough to make the required operations (essentially, computing the common tangents of two homothetic copies of $\mathscr{B}$) in constant time.

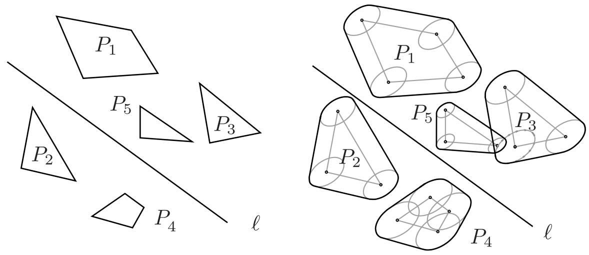

Fig. 3. Decision problem interpreted via Minkowski sums.

We can solve the decision problem applying the same technique that was used to solve the Euclidean version of the problem. We use duality and observe that the set of lines intersecting a convex set $P_i \oplus \omega_i\delta\mathscr{B}$ corresponds in dual plane to the region between two convex curves, U_i and L_i, which are dual to the set of upper and lower tangents of $P_i \oplus \omega_i\delta\mathscr{B}$, respectively. Therefore, the set of lines which are at weighted distance at least δ from $P_1, P_2, \ldots, P_m$ corresponds to a set of faces in the arrangement formed by those curves. Using exactly the same arguments as in Lemma 2 it can be shown that the arrangement formed by the curves U_i and L_i has complexity $O(m^2 + n)$ and can be computed in time $O((m^2 + n\log m)\log n)$. Therefore, we have:

Theorem 3. *Let $\mathscr{P} = \{P_1, \ldots, P_m\}$ be a set of convex polygons with a total of n vertices and let $\omega_1, \ldots, \omega_m$ be the corresponding weights. Let $d_{\mathscr{B}}$ be the distance defined by the unit ball $\mathscr{B}$. The problem [Dec] can be solved in time $O((m^2 + n\log m)\log n)$ using $O(m^2 + n)$ space.*

4.1. The optimization problem

Parametric search is a well known technique in geometric optimization which originated in [24]. It can be used to solve optimization problems which are *monotone* with respect to a given parameter δ: if the answer to the corresponding decision problem is positive for a given δ_1, then it is also positive for every $\delta_2 < \delta_2$. Therefore, if we denote by δ^* the value of the parameter associated to the optimal solution, we know that if the answer to the decision problem is positive for δ_1 then $\delta^* \geqslant \delta_1$, while if it is negative, then $\delta^* < \delta_1$. The interested reader can find both theoretically oriented and applied oriented expositions of the parametric search technique in [2,31].

[4] If the number of edges is not a constant, the same algorithm works, but the complexity is related to the number of edges of $\mathscr{B}$.

Perhaps the main difficulty of the parametric search technique is that we need a parallel algorithm for the generic version of the decision problem, which is not always easy. Nevertheless, it can be observed that it is not necessary that the generic algorithm solves the problem under consideration: all we need is that the output of the generic algorithm changes combinatorially at δ^*. Actually, in quite some cases sorting can play the role of the generic algorithm and, in these cases, we can use one of the parallel sorting algorithms. It has been pointed out in [28] that Cole's algorithm presented in [6] may be specially appropriated in practice because it has good asymptotic complexity and small constants, and that quicksort can also give good results in practice. Let us see why sorting can also be used for our problem.

Let us consider $P_i^\delta = P_i \oplus \delta\omega_i\mathscr{B}$. The (curved) polygon P_i^δ is made up of arcs of copies of the unit ball $\mathscr{B}$ centered at vertices of P_i and tangents between them. For a vertex $u \in P_i$, we denote by u^δ the corresponding arc of P_i^δ. As δ increases, the two inner tangents to P_i^δ and P_j^δ rotate in opposite directions. We denote by $\tau_{ij}^+(\delta)$ the inner tangent to P_i^δ and P_J^δ that rotates counterclockwise and by $\tau_{ij}^-(\delta)$ the inner tangent that rotates clockwise. Finally, $\alpha_{ij}^+(\delta)$ and $\alpha_{ij}^-(\delta)$ are, respectively, the angles determined by the inner tangents with the OX axis (see Fig. 4a).

The key observation is that, if we sort the $O(m^2)$ angles $\alpha_{ij}^+(\delta)$ and $\alpha_{ij}^-(\delta)$, the order changes in the candidate solutions to the maximin line problem. Actually, for candidates of type (a) we have that $\alpha_{ij}^+(\delta^*) = \alpha_{ij}^-(\delta^*)$ while for candidates of type (b) it holds $\alpha_{ij}^+(\delta^*) = \alpha_{ij}^-(\delta^*)$ for a suitable labeling of the polygons (see Fig. 4b).

Our next step is computing the set of inner tangents between every pair of polygons at the optimal solution, $\tau_{ij}^\pm(\delta^*)$.

Lemma 3. *The set* $\tau_{ij}^\pm(\delta^*)$ *can be computed in time* $O(nm + (m^2 + n\log m)\log^2 n)$.

Proof. The inner tangents to P_i and P_j, i.e., $\tau_{ij}^\pm(0)$, can be computed in time $O(\log n_i + \log n_j)$ (see [30]). Let us assume that we label the vertices of P_i and P_j counterclockwise and in such a way that $\tau_{ij}^+(0)$ is tangent to P_i and P_j at vertices u_1 and v_1, respectively (see Fig. 5). Now, as δ increases, $\tau_{ij}^+(\delta)$ is tangent to two copies of the unit ball $\mathscr{B}$ centered at u_1 and v_1 and scaled according with ω_i and ω_j, the weights corresponding to the polygons. We recall that, in our model of computation, this set of tangents can be computed and described in constant time. The vertices supporting the inner tangent change either when the tangent to P_i^δ at the edge $u_1^\delta v_2^\delta$ is also tangent to P_i^δ (at u_1^δ) or when the tangent to P_i^δ at the edge $u_1^\delta u_2^\delta$ is also tangent to P_j^δ (at v_1^δ).

This process ends when we reach a value of δ for which polygons P_i^δ and P_j^δ are tangent. Because we always move forward on the boundary of the polygons, namely, counterclockwise for τ_{ij}^+ and clockwise for τ_{ij}^-, the set of values of δ for which the vertices supporting the tangents change has size at most $n_i + n_j$.

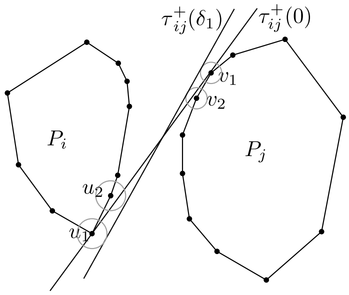

Fig. 5. Computing $\tau_{ij}^+(\delta)$.

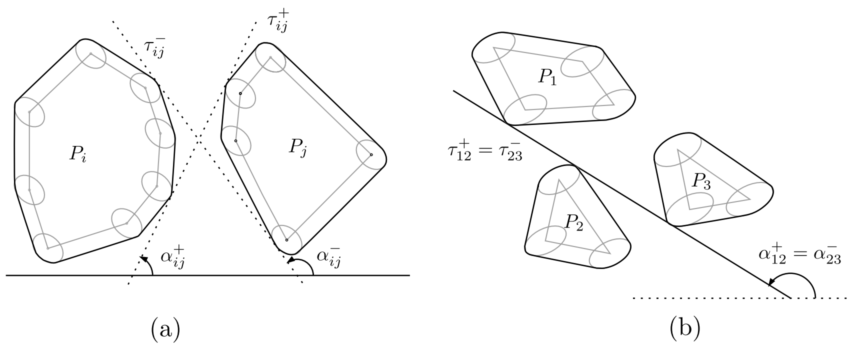

Fig. 4. Reduction to sorting inner tangents.

Repeating this process for every pair of polygons, we get a set of $\sum_{i\neq j}(n_i+n_j)=\mathrm{O}(nm)$ values of δ where the vertices supporting a member of $\tau^{\pm}_{ij}(\delta)$ change. Now, we can use the solution for the decision problem and the monotonicity of the optimization problem in order to locate the interval corresponding to δ^*: we compute the median δ_m and run the decision algorithm for δ_m in order to decide whether $\delta^*<\delta_m$ or $\delta^*\geqslant\delta_m$. By iterating this process we find the interval containing δ^*. For the complexity, observe that iterative median finding can be done in time $\mathrm{O}(nm)$ and that we run the sequential algorithm $\mathrm{O}(\log n)$ times. After elementary operations, we obtain the claimed complexity. □

Now, we can order the $\mathrm{O}(m^2)$ inner tangents. Observe that, although we have not computed δ^*, we have been able to compute the arcs which support the inner tangents at the critical value. Therefore, in the following we do not have to deal with polygons but only with *elementary arcs*: given a vertex $u_i\in P_i$, the elementary arc corresponding to u_i is the arc of the ball $u_i^{\delta}=u_i\oplus\delta\omega_i\mathscr{B}$ in the boundary of $P_i\oplus\delta\omega_i\mathscr{B}$.

We need to study the set of solutions of the equation $\alpha^{\pm}_{ij}(\delta)=\alpha^{\pm}_{kl}(\delta)$. Clearly, $\alpha^{+}_{ij}(\delta)$ and $\alpha^{-}_{kl}(\delta)$ intersect in at most one point, because the former one is an increasing function while the later is a decreasing function, and the same is true for $\alpha^{-}_{ij}(\delta)$ and $\alpha^{+}_{kl}(\delta)$. In the next result we study the set of solutions of the equation $\alpha^{+}_{ij}(\delta)=\alpha^{+}_{kl}(\delta)$ (the situation for the clockwise rotating tangents is identical).

Lemma 4. *Let $\alpha^{+}_{ij}(\delta)$ and $\alpha^{+}_{kl}(\delta)$ be elementary arcs describing the counterclockwise rotating inner tangents to u_i^{δ} and u_j^{δ} and to u_k^{δ} and u_l^{δ}, respectively. Let us assume that at least one of the indices k, l is different from i and j. If the equation $\alpha^{+}_{ij}(\delta)=\alpha^{+}_{kl}(\delta)$ has more than one solution, then the curves $\alpha^{+}_{ij}(\delta)$ and $\alpha^{+}_{kl}(\delta)$ are equal whenever both are defined.*

Proof. Without loss of generality, we can assume that k is different from i and j. It is well known (see for instance [15]) that all inner tangents to balls u_i^{δ} and u_j^{δ} pass through a fix point, that we denote m_{ij}. This point is on the line defined by u_i and u_j and it holds that $\frac{\mathrm{d}(u_i,m_{ij})}{\mathrm{d}(u_j,m_{ij})}=\frac{\omega_i}{\omega_j}$.

Let us assume that α^{+}_{ij} and α^{+}_{kl} intersect twice. In Fig. 6 the angles for which the common tangents are parallel are 0 and γ. Let α_h, β_h, for $h=1,2$, be the tangency points with balls u_i^{δ} and u_k^{δ}. We observe that the triangles $\alpha_1\alpha_2u_i$ and $\beta_1\beta_2u_k$ are similar and, therefore, the segments $\alpha_1\alpha_2$ and $\beta_1\beta_2$ are parallel and such that $\frac{\mathrm{d}(\alpha_1,\alpha_2)}{\mathrm{d}(\beta_1,\beta_2)}=\frac{\omega_i}{\omega_k}$. Because the triangles $\alpha_1\alpha_2m_{ij}$ and $\beta_1\beta_2m_{kl}$ are also similar, it follows that $\frac{\mathrm{d}(\alpha_1,m_{ij})}{\mathrm{d}(\beta_1,m_{kl})}=\frac{\omega_i}{\omega_k}$ and, therefore, triangles $u_i\alpha_1m_{ij}$ and $\beta_1\beta_2m_{kl}$ are similar too. But then the edges u_im_{ij} and u_km_{kl} are parallel and their lengths are in the proportion $\frac{\omega_i}{\omega_k}$, which implies that curves $\alpha^{+}_{ij}(\delta)$ and $\alpha^{+}_{kl}(\delta)$ are equal whenever both are defined. □

We are now ready to describe the final procedure: we order the $\mathrm{O}(m^2)$ inner tangents in $\mathrm{O}(\log m)$ steps, and perform $\mathrm{O}(m^2)$ comparisons in each step. According to our model of computation and using Lemma 4, we can guarantee that each comparison can be performed in constant time and has at most one solution. Therefore, the complexity of the final sorting step of the algorithm is $\mathrm{O}(m^2\log m+T_{\mathrm{s}}\log^2 m)$, where $T_{\mathrm{s}}=\mathrm{O}((m^2+n\log m)\log n)$. From Lemma 3, we know that the complexity of computing the set of inner tangents at δ^* is $\mathrm{O}(nm+(m^2+n\log m)\log^2 n)$. Putting all this together, we have:

Theorem 4. *Let $\mathscr{P}=\{P_1,\ldots,P_m\}$ be a set of convex polygons with a total of n vertices and let ω_i be the corresponding weights. The problem [Opt] can be solved in time $\mathrm{O}(nm+n\log^2 n\log m+m^2\log n\log^2 m)$ using $\mathrm{O}(m^2+n)$ space.*

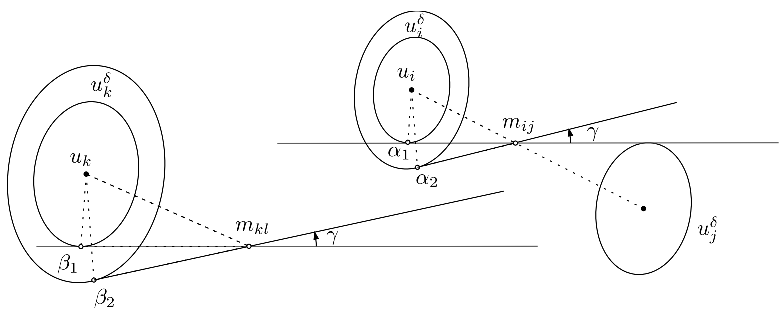

Fig. 6. Illustration for the Proof of Lemma 4.

It may be argued that the algorithm is too involved and that the complexity is too high in order to be useful in practice. Nevertheless, we observe that the third term in the complexity bound only dominates if the number of polygons is close to the number of vertices. On the other hand, if m is small compared to total number of vertices, namely, if $m = O(n^{1-\varepsilon})$ for some $0 < \varepsilon < 1$, then the complexity of the algorithm is $O(n^{2-\varepsilon})$. Finally, if we have only a constant number of polygons with a total of n vertices, then the complexity is $O(n \log^2 n)$.

5. The three-dimensional scene

In this section we deal with the problem of computing the plane which maximizes the minimum distance to a set of polyhedra in $\mathbb{R}^3$.

Given a set $\mathscr{P} = \{P_1, \ldots, P_m\}$ of m polyhedra in $\mathbb{R}^3$ with a total of n vertices, we want to find a plane π such that

- $\pi \cup CH(\mathscr{P}) \neq \emptyset$.
- $\min_i d(P_i, \pi)$ is maximum.

This problem is named as the *obnoxious plane problem* in [9], where it is solved in $O(n^3)$ time and $O(n^2)$ space. The problem is equivalent to finding an empty region bounded by two parallel planes as wide as possible and defining a non-trivial partition in the set of polyhedra, that we refer as *the widest empty slab problem*. Let S be the set of vertices of the objects in $\mathscr{P}$.

We state a necessary condition for slab optimality in arbitrary dimension. We say that a hyperplane π strictly separates two sets of points if each of the sets is contained in one of the open halfspaces defined by π.

Theorem 5. *Let π^* be a solution to an instance of the obnoxious hyperplane problem and let π_1 and π_2 be the bounding hyperplanes of the slab generated by π^*. Then, the sets $S_1 = S \cap \pi_1$ and $S_2 = S \cap \pi_2$ cannot be strictly separated by a hyperplane orthogonal to π^*.*

Proof. Let us assume that $S_1 = \{p_i\}_{i \in I}$ and $S_2 = \{q_j\}_{j \in J}$. We denote by $\vec{n}$ the unitary vector orthogonal to π^*. Then, the distance between the parallel planes π_1 and π_2 is $\Delta = |\vec{n} \cdot p_i\vec{q}_j|$. First, observe that if S_1 and S_2 can be strictly separated by a hyperplane h orthogonal to π^*, then the unitary vector normal to h, denoted by $\vec{v}$, can be chosen such that

$$\min_{p_i \in S_1, q_j \in S_2} \vec{v} \cdot p_i\vec{q}_j = k > 0.$$

Now we consider an empty slab orthogonal to $\vec{n}_\in = \vec{n} + \varepsilon\vec{v}$. The width of the slab is

$$\Delta(\varepsilon) = \min_{p_i \in S_1, q_j \in S_2} \frac{\vec{n}_\in \cdot p_i\vec{q}_j}{\|\vec{n}_\varepsilon\|} = \min_{p_i \in S_1, q_j \in S_2} \frac{\Delta + \varepsilon\vec{v} \cdot p_i\vec{q}_j}{1 + \varepsilon^2} = \frac{\Delta + k\varepsilon}{1 + \varepsilon^2}.$$

Because $\Delta(0) = \Delta$ and $\Delta'(0) = k > 0$, we can guarantee that $\Delta(\varepsilon) > \Delta$ for $\varepsilon > 0$ small enough. □

As a consequence of the preceding theorem we can restrict our search to slabs C that satisfy one of the four following conditions that are illustrated in Fig. 7 for the case of dimension 3. See [9] for an alternative proof.

Following the same approach as in the two dimensional case, we solve the problem by exploring efficiently only the set of planes avoiding $\mathscr{P}$. We first give a bound for the combinatorial complexity of this set and show how it can be computed.

Let $\mathscr{D}$ be the transformation which maps a point $p = (a, b, c)$ to the plane $\mathscr{D}(p) : z = ax + by - c$ in the dual space, and maps a non-vertical plane π: $z = mx + ny - d$ to the point $\mathscr{D}(\pi) = (m, n, d)$ in the dual space. Given a polyhedron P_i, we denote by A_i the set of planes avoiding P_i and by A_i^* the set of points dual to planes in A_i.

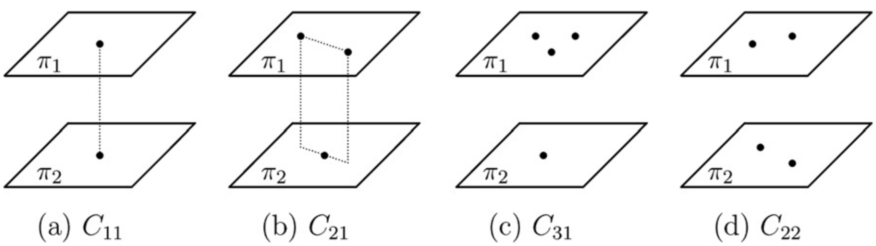

Fig. 7. Types of candidate slabs according to Theorem 5.

Theorem 6. *Let $\mathscr{P} = \{P_1, \ldots, P_m\}$ be a set of m polyhedra with a total of n vertices. The set of planes avoiding $\mathscr{P}$ has complexity $O(m^2n)$ and can be computed in time $O(m^2n\log(m^2n))$.*

Proof. We want to argue that the number of vertices of $\mathscr{A} = \cap_{i=1}^m A_i^*$ is $O(m^2n)$. The vertices of $\mathscr{A}$ correspond to planes passing through three vertices of $\mathscr{P}$ and can be classified into three different types:

1. planes containing three vertices of a polyhedron (therefore, containing a face of its convex hull),
2. planes passing through an edge of a polyhedron and a vertex of another polyhedron,
3. planes passing through three vertices of three different polyhedra.

The number of vertices of the first type is clearly $O(n)$ and, for the second type, we observe that there are at most two planes tangent to a given edge and a given polyhedron and, therefore, the number of those vertices is $O(mn)$. Finally, for the third type of vertices, we claim that the number of planes tangent to polyhedra P_i, P_j and P_k, with n_i, n_j and n_k vertices respectively, is $O(n_i + n_j + n_k)$. In order to prove the claim, observe that A_i^* is the union of two unbounded convex polyhedra. Therefore, $A_i^* \cap A_j^* \cap A_k^*$ can be described as the union of at most eight disjoint sets, each of which is the intersection of three convex polyhedra and has complexity $O(n_i + n_j + n_k)$. The first part of the proof is finished because

$$\sum_{i \neq j \neq k} O(n_i + n_j + n_k) = O(m^2n).$$

In order to compute $\mathscr{A}$ we use a divide and conquer approach. Assume that we partition $\mathscr{P}$ into two sets of $\lceil \frac{m}{2} \rceil$ and $\lfloor \frac{m}{2} \rfloor$ polyhedra, denoted $\mathscr{R}$ and $\mathscr{B}$, with n vertices in total. Let $\mathscr{A}_r$ and $\mathscr{A}_b$ denote, respectively, the sets of points dual to planes avoiding $\mathscr{R}$ and $\mathscr{B}$, respectively. The crux of the method is observing that the merge step reduces to computing the intersection of $\mathscr{A}_r = \cup_{i=1}^k R_i$ and $\mathscr{A}_b = \cup_{i=1}^l R_i$, where R_i and B_j are convex polyhedra and the total complexity of $\mathscr{A}_r$ and $\mathscr{A}_b$ is $O(m^2n)$.

We compute the intersection of $\mathscr{R}$ and $\mathscr{B}$ using a space sweeping approach. It is clear that the intersection can be easily computed if we are able to maintain the planar subdivision generated in the sweep plane by one of the sets, say $\mathscr{R}$, and perform point location in such subdivision when we encounter a new vertex of $\mathscr{B}$. These operations can be done efficiently by using the dynamic point location structure of Goodrich and Tamassia [14] which can manage monotone subdivisions (in our case, the subdivision is convex) and takes $O(\log n)$ per update and $O(\log^2 n)$ per point location query, where n is the total size of the subdivision. Because the total size of $\mathscr{R}$ and $\mathscr{B}$ is $O(m^2n)$, it follows that the intersection can be computed in time $O(m^2n\log(m^2n))$. Therefore, if we denote by $T(m,n)$ the time required by the whole algorithm, we obtain the recursive formula

$$T(m,n) = T(m/2, n_1) + T(m/2, n - n_1) + O(m^2n\log(m^2n)),$$

which solves to $T(m,n) = O(m^2n\log(m^2n))$. □

We now describe how to use the arrangement $\mathscr{A} = \cap_{i=1}^m A_i^*$ in order to solve the problem. As shown is the two dimensional case, the idea is to solve the optimization problem visiting all cells c in $\mathscr{A}$ and identifying the candidate slabs associated with c. By using the properties of the duality transform we look at open vertical segments whose endpoints lie on the boundary of each cell. We have to examine all the vertical segments inside a cell that correspond with candidates of type C_{11}, C_{21}, C_{31}, C_{22} as illustrated in Fig. 7.

When leaving a cell c, we test every face–face, edge–face, vertex–face and edge–edge pair of c in order to identify and compute the width of all pairs that are vertically aligned, i.e., the widths of the candidate slabs in the primal space.

A detailed description of the detection of candidates within a cell in a three-dimensional arrangement is given in [9]. In each cell, each candidate can be processed in $O(1)$ amortized time. At this point, we should note that candidates type C_{11} and C_{21} differ from candidates C_{31} and C_{22}. In fact, the number of vertical segments associated with a face–face or edge–face pair is not finite. However, the orthogonality condition of the former can be used to identify those types of candidates in amortized $O(1)$ time per cell.

As a consequence of the above description, the overall time we need to obtain all the candidates in our arrangement and compute the optimal one is proportional to the size of the arrangement, and we have established the following result:

Theorem 7. *An obnoxious plane through a set of m polyhedral objects in $\mathbb{R}^3$ with a total of n vertices can be computed in $O(m^2n\log(m^2n))$ time and $O(m^2n)$ space.*

Acknowledgements

The authors wish to thank Boris Aronov for his help in the Proof of Theorem 6 and an anonymous referee whose comments suggested an improvement in the solution of problem [Opt] in Section 4.

References

[1] P.K. Agarwal, M. Sharir, Ray shooting amidst polygons in 2D, Journal of Algorithms 21 (1996) 508–519.
[2] P.K. Agarwal, M. Sharir, Efficient algorithms for geometric optimization, ACM Computing Surveys 30 (1998) 412–458.
[3] J. Brimbeg, H. Juel, A. Schöbel, Linear facility location in three dimensions – models and solution methods, Operations Research 50 (6) (2002) 1050–1057.
[4] J. Brimberg, G.O. Wesolowsky, Locating facilities by minimax relative to closest points of demand areas, Computers & Operations Research 29 (2002) 625–636.
[5] E. Carrizosa, M. Muñoz-Márquez, J. Puerto, The Weber problem with regional demand, European Journal of Operational Research 104 (2) (1998) 358–365.
[6] R. Cole, Parallel merge sort, SIAM Journal on Computing 17 (1988) 770–785.
[7] J.M. Díaz-Báñez, F. Gómez, G. Toussaint, Computing shortest paths for transportation of hazardous materials in continuous spaces, Journal of Food Engineering 70 (2005) 293–298.
[8] J.M. Díaz-Báñez, F. Hurtado, Computing obnoxious 1-corner polygonal chains, Computers & Operations Research 33 (2005) 1117–1128.
[9] J.M. Díaz-Báñez, M. López, J.A. Sellarès, Locating an obnoxious plane, European Journal of Operational Research 173 (2006) 556–564.
[10] J.M. Díaz-Báñez, J.A. Mesa, A. Shöbel, Continuous location of dimensional structures, European Journal of Operational Research 152 (2004) 22–44.
[11] Z. Drezner, in: Zvi Drezner (Ed.), Facility Location: A Survey of Applications and Methods, Springer, 1995.
[12] Z. Drezner, G.O. Wesolowsky, Location of an obnoxious route, Journal of Operational Research Society 40 (1989) 1011–1018.
[13] H. Edelsbrunner, L. Guibas, Topologically sweeping an arrangement, Journal of Computer and System Sciences 38 (1989) 165–194.
[14] M.T. Goodrich, R. Tamassia, Dynamic trees and dynamic point location, SIAM Journal on Computing 28 (1998) 612–636.
[15] G.H. Hardy, A Course of Pure Mathematics, Cambridge University Press, 1967, pp. 99–106.
[16] H.W. Hamacher, S. Nickel, Multicriteria planar location problems, European Journal Of Operational Research 94 (1996) 66–86.
[17] Y. Hinojosa, Some problems in location theory, Ph.D. Thesis, University of Seville, 2000.
[18] Y. Hinojosa, J. Puerto, The polyhedral norm approach to the problem of locating obnoxious routes, Studies in Locational Analysis 12 (1999) 49–65.
[19] M. Houle, A. Maciel, Finding the widest empty corridor through a set of points, in: G.T. Toussaint (Ed.), Snapshots of Computational and Discrete Geometry, TR SOCS-88.11, Department of Computer Science, McGill University, Montreal, Canada, 1988, pp. 210–213.
[20] M. Houle, G.T. Toussaint, Computing the width of a set, IEEE Transactions on Pattern Analysis and Machine Intelligence 10 (1988) 760–765.
[21] J. Janardan, F.P. Preparata, Widest-corridor problems, Nordic Journal of Computing 1 (1994) 231–245.
[22] N.M. Korneenko, H. Martini, Hyperplane approximation and related topics, in: János Pach (Ed.), New Trends in Discrete and Computational Geometry, Springer-Verlag, New York, 1993, pp. 135–162, Chapter 6.
[23] D.T. Lee, Y.F. Wu, Geometric complexity of some location problems, Algorithmica 1 (1986) 193–211.
[24] N. Megiddo, Applying parallel computation algorithms in the design of serial algorithms, Journal of ACM 30 (1983) 852–865.
[25] J.G. Morris, J.P. Norback, A simple approach to linear facility location, Transportation Science 14 (1) (1980) 1–8.
[26] J.G. Morris, J.P. Norback, Linear facility location – solving extensions on the basic problems, European Journal of Operational Research 12 (1983) 90–94.
[27] S. Nickel, J. Puerto, A.M. Rodriguez-Chía, An approach to location models involving sets as existing facilities, Mathematics of Operational Research 28 (2003) 693–715.
[28] R.V. Oostrum, R.C. Veltkamp, Parametric search made practical, Computational Geometry 28 (2004) 75–88.
[29] F. Plastria, Continuous covering location problems, in: Zvi Drezner, Horst W. Hamacher (Eds.), Facility Location: Applications and Theory, Springer, 2002.
[30] H. Rohnert, Shortest paths in the plane with convex polygonal obstacles, Information Processing Letters (23) (1986) 71–76.
[31] J. Salowe, Parametric search, in: J.E. Goodman, J. O'Rourke (Eds.), Handbook of Discrete and Computational Geometry, 1997, pp. 683–698.
[32] M. Sharir, P.K. Agarwal, Davenport-Schinzel Sequences and Their Geometric Applications, Cambridge-University Press, 1995.
[33] A. Shöbel, Locating least-distant lines in the plane, European Journal Of Operational Research 106 (1998) 152–159.