A Near-Linear Algorithm for the Planar Segment Center Problem[*]

Alon Efrat[†]   Micha Sharir[‡]

May 14, 1996

Abstract

Let P be a set of n points in the plane and let e be a segment of fixed length. The segment center problem is to find a placement of e (allowing translation and rotation) which minimizes the maximum euclidean distance from e to the points of P. We present an algorithm that solves the problem in time $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, improving the previous solution of Agarwal et al. [3] by nearly a factor of $O(n)$.

1   Introduction

The problem considered in this paper is: "Given a set S of n points in the plane and a segment e, find a *center placement* of e (allowing translations and rotations) at which the maximum (euclidean) distance from e to the points of S is minimized (see Figure 1); the distance between e and a point p is the minimum euclidean distance $\min_{q \in e} |pq|$."
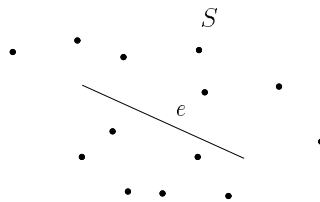

Figure 1: The segment center problem

The problem was posed by Imai et al. [17] a few years ago. It generalizes the well known notions of the point center (which is the center of the smallest disk containing S) [23] and of the line-center [18]. Finding the point center and the line center of a set S is easy, and can be done in time $O(n)$ and $O(n \log n)$, respectively. The segment-center problem appears to be more difficult. An initial and rather inefficient solution to the problem was given by Imai et al. [17]; it runs in time $O(n^4 \log n)$. A considerably improved solution was later given by Agarwal et al. [3]; their solution

uses the parametric searching technique of Megiddo [23], and requires $O(n^2 \alpha(n) \log^3 n)$ time, where $\alpha(n)$ is the inverse Ackermann function.

In this paper we further improve the solution, reducing the running time to $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε. Our solution is also based on the parametric searching technique, and we assume familiarity of the reader with that technique (details can be found, e.g., in the paper [3] just cited). The parametric searching paradigm suggests that we first solve the 'fixed-size' problem: Given S and e as above and a real parameter $d > 0$, determine whether there exists a placement of e at which all points of S lie within distance d from e. Without loss of generality, assume that the length of e is 1. To solve this problem, we define $H(e, d) = e \oplus B_d$ to be an expansion of e by distance d (here B_d is the disk of radius d about the origin, and $e \oplus B_d$ denotes the *Minkowski sum* of e, at some standard placement, and B_d). $H(e, d)$ has the shape of a *hippodrome* — a rectangle of dimensions $1 \times 2d$ with two semicircles of radius d attached to its sides. Next, instead of moving e about, we fix e (and $H(e, d)$) and move S rigidly. It is easily seen that the problem reduces to that of determining whether $P \equiv conv(S)$ can be placed (by translations and rotations) inside $H(e, d)$; see Figure 2.
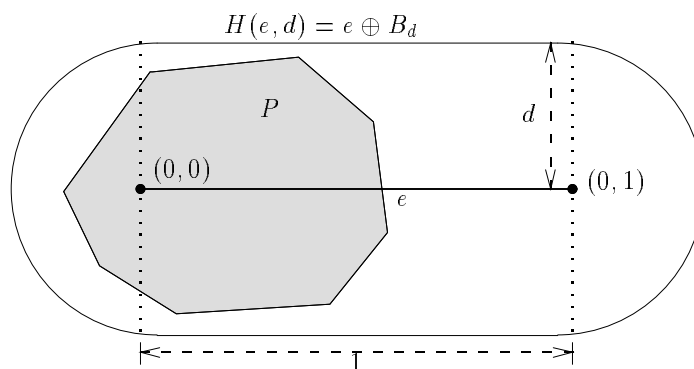

Figure 2: The transformed segment center problem

The problem has thus been reduced to a *polygon containment* problem, such as those studied in [7, 19, 25], with the difference that the environment $H(e, d)$ in which P has to be placed is not polygonal, although its shape is very simple. Nevertheless, techniques similar to those used in [19, 25] can be developed. That is, one can easily show that there exists a free placement of P inside $H(e, d)$ iff there exists a *critical free placement*, at which $P \subseteq H(e, d)$ and three vertices of P touch $\partial H(e, d)$ (see [3] for details). Thus our goal is to determine whether there exists a critical free placement of P. We accomplish this using a somewhat indirect method, which exploits additional geometric structure of the problem.

We derive sequential and parallel versions of the algorithm. The sequential algorithm runs in time $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, and the parallel algorithm runs (in Valiant's comparison model [28]) in $O(\log^2 n)$ time and uses $O(n^{1+\varepsilon})$ processors.

The improvement that we achieve is a consequence of an improved combinatorial bound on the number of critical placements of the polygon P in $H(e, d)$. Specifically, we show that the number of such placements is only $O(n \log n)$. This improved bound is obtained by applying a more careful geometric analysis of the structure of such placements, which exploits, among other things, a combinatorial bound of Capoyleas and Pach [6] concerning diagonals in a convex polygon, and the results of Bienstock and Györi [5] and of Füredi [15] concerning 0-1 matrices with a certain forbidden submatrix configuration.

After establishing these improved combinatorial bounds in Section 2, we present in Section 3 an

efficient algorithm for determining whether P can be placed inside $H(e, d)$, with the performance bounds reported above. The design and analysis of the algorithm exploit the near-linear number of the critical placements, The resulting sequential algorithm is relatively easy to parallelize (in Valiant's model). We briefly describe the parallel version in Section 4, and then plug all this into the parametric searching machinery to obtain our main result:

Theorem 1.1 *Given a set S of n points in the plane, and a line segment e, we can compute a center placement for e, which minimizes the maximum distance from e to the points of S, in time $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$.*

We remark that our technique can also be used to solve the following related problem: Given a set $\mathcal{D}$ of n congruent discs, find the shortest segment e that meets all the discs of $\mathcal{D}$. The algorithm given below can be adapted to solve this problem in time $O(n^{1+\varepsilon})$. See [4] for some related results.

2 Combinatorial Bounds

Before tackling the analysis of critical placements within a hippodrome, we first study the following simpler problem.

2.1 Critical contacts with a semicircle

Let P be a convex polygon whose vertices are $p_1, \ldots, p_n$, appearing in this order in clockwise direction along ∂P. Let Γ be a semicircle of some fixed radius d. We assume that P lies in general position with respect to d, meaning that no circle of radius d passes through three distinct vertices of P. We define an ordered pair (p_i, p_j) of vertices of P to be *conjugate* if Γ can be placed so that it passes through p_i and p_j, so that p_j lies clockwise from p_i along Γ, and so that Γ avoids the interior of P. Our goal is to obtain a sharp upper bound on the number of conjugate pairs among the vertices of P.

For any two points x, y in the plane, we define their *arc-distance* $\sigma(x, y)$ to be the length of the (smaller) circular arc of radius d that connects x to y. That is,

$$\sigma(x, y) = 2d \arcsin \frac{|xy|}{2d} \, ,$$

where $|xy|$ is the euclidean distance between x and y; if $|xy| > 2d$, we put $\sigma(x, y) = +\infty$.

Lemma 2.1 *Let a, b, c be three vertices of P appearing clockwise along ∂P in this order, and assume that (a, c) is a conjugate pair. Let Γ_1 be a copy of Γ which passes through b, avoids a and the interior of P, and suppose that the ray emerging from a towards c intersects Γ_1 from inside its disk D_1. Then the length of the counterclockwise portion Γ_1^* of Γ_1 delimited by b is smaller than $\sigma(a, b)$.*

Proof: Refer to Figure 3. Let Γ_0 be the copy of Γ that passes through a and c and avoids the interior of P, and let D_0 be the disk of Γ_0. We claim that Γ_0 and Γ_1 intersect at most once. Indeed, if they intersected twice, then, as is easily seen, the disk D_1 would have to be disjoint from the interior of P, so the ray from a to c could not hit Γ_1 from inside the disk D_1, contrary to assumption. Moreover, Γ_0 and Γ_1 have to intersect once, because $b \in \Gamma_1$ lies strictly inside D_0 (because of the general position assumption), whereas the point c' of intersection between Γ_1 and the ray from a to c lies outside D_0 (or on its boundary). In particular, this argument implies that the portion Γ_1^* must be fully contained in the interior of the disk D_0.

Moreover, the counterclockwise endpoint x of Γ_1^* and a must lie on the same side of the line ℓ passing through b and perpendicular to ac. This follows from the observation that the segment of ℓ
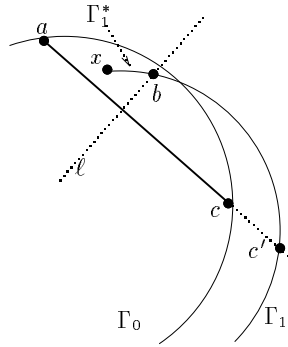
Figure 3: The configuration in the proof of Lemma 2.1

between b and its intersection with ac is contained in P, and thus also in the interior of D_1, which easily implies that Γ_1 must continue counterclockwise from b into the appropriate side of ℓ. Let D_0' be the portion of D_0 delimited by ac and by ℓ and containing x. Using elementary geometric arguments, one easily verifies that $|bx| < |ab|$. Since $\sigma(x, y)$ is a strictly increasing function of $|xy|$, we obtain

$$|\Gamma_1^*| = \sigma(b, x) < \sigma(a, b) \ ,$$

as asserted. $\square$

Corollary 2.2 *Let a, b, c be three vertices of P appearing clockwise along ∂P in this order, and assume that (a, c) and (b, c) are both conjugate pairs. Let Γ_1 be a copy of Γ that passes through b and c and avoids a and the interior of P. Then the length of the portion of Γ_1 between its counterclockwise endpoint and b is smaller than $\sigma(a, b)$.*

Proof: The conditions of the previous lemma obviously hold in this case. $\square$

Lemma 2.3 *Let a, b, c be three vertices of P appearing clockwise along ∂P in this order, and assume that (a, c) is a conjugate pair. Then*

$$\sigma(a, b) + \sigma(b, c) < \sigma(a, c) \ .$$

Proof: Refer to Figure 4. Let Γ_0 be the copy of Γ that passes through a and c and avoids the interior of P, and let D_0 be the disk of Γ_0. Draw the perpendicular line ℓ to ac through b and extend it until it meets Γ_0 at a point b' (b itself clearly lies inside D_0). It is easily checked that $|ab| < |ab'|$ and $|bc| < |b'c|$. Since $\sigma(x, y)$ is an increasing function of $|xy|$, we have also $\sigma(a, b) < \sigma(a, b')$ and $\sigma(b, c) < \sigma(b', c)$. Since $\sigma(a, b') + \sigma(b', c) = \sigma(a, c)$, the lemma follows. $\square$

We can now prove the main result of this section:

Theorem 2.4 *Given a convex n-gon P and a fixed radius d in general position, the maximum number of conjugate pairs of vertices of P, with respect to a semicircle of radius d, is $O(n \log n)$.*

Proof: Call a conjugate pair (a, b) *short* if $\sigma(a, b) \leq \frac{\pi}{3} d$, and *long* otherwise.

Claim I: The number of short conjugate pairs is $O(n)$.

Proof: (We are grateful to János Pach for a helpful discussion that has led to this proof.) We show that there do not exist three short conjugate pairs (a, u), (b, v) and (c, w) so that a, b, c, u, v, w appear in this clockwise order along ∂P. By a recent result of Capoyleas and Pach [6], if $\mathcal{C}$ is a
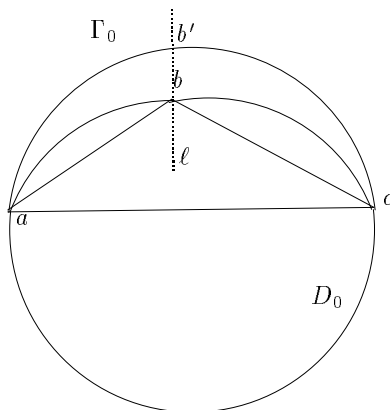
Figure 4: The configuration in the proof of Lemma 2.3

set of diagonals of a convex polygon with n vertices, such that no three distinct diagonals in $\mathcal{C}$ are pairwise intersecting, then the size of $\mathcal{C}$ is at most $6n$. This result thus implies that the number of short conjugate pairs is at most $6n$. Suppose to the contrary that there exist three short conjugate pairs as above. Let Γ_0 denote a copy of Γ that passes through b and v and avoids the interior of P; let x and y denote respectively the counterclockwise and clockwise endpoints of Γ_0; see Figure 5.
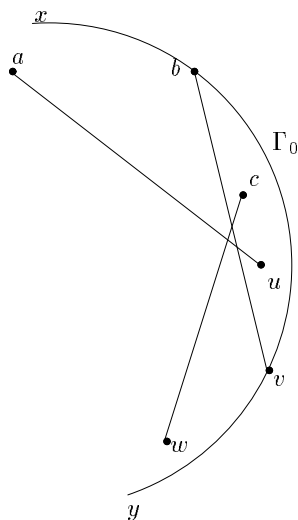
Figure 5: Three short pairwise interleaving conjugate pairs

Applying Lemma 2.1 to each of the pairs (a, u) and (c, w) (where in the second application we interchange the role of clockwise and counterclockwise directions), we conclude that $\sigma(x, b) < \sigma(a, b)$ and $\sigma(v, y) < \sigma(v, w)$. Hence,

$$\pi d = \sigma(x, b) + \sigma(b, v) + \sigma(v, y) < \sigma(a, b) + \sigma(b, v) + \sigma(v, w) .$$

Applying Lemma 2.3 to each of the pairs (a, u) and (c, w), we conclude that

$$\sigma(a, b) < \sigma(a, u)   \text{and}   \sigma(v, w) < \sigma(c, w) .$$

Hence, since all these pairs are short,

$$\pi d < \sigma(a, u) + \sigma(b, v) + \sigma(c, w) \le \frac{\pi}{3}d + \frac{\pi}{3}d + \frac{\pi}{3}d = \pi d ,$$

a contradiction which establishes the claim. $\Box$

We next analyze the number of long conjugate pairs. For this we partition ∂P into a constant number of connected portions, so that the turning angle along each portion is less that $\frac{\pi}{3}$. More precisely, each portion starts and ends at the midpoints of two edges of P, so that the angle between the outward normals of these two edges is less than $\frac{\pi}{3}$. If P has a vertex v whose interior angle is $< \frac{2\pi}{3}$, we also break ∂P at v (there can be only a constant number of such vertices v).

Claim II: If (a, b) is a long conjugate pair, the angle between the outward normals of the edge preceding a (i.e., incident to a and lying counterclockwise to it) and the edge succeeding b (incident to b and lying clockwise to it) is greater than $\frac{\pi}{3}$.

Proof: Let Γ_0 be a copy of Γ that passes through a and b and avoids the interior of P. The angle in the claim is larger than the angle between the two outward normals to Γ_0 at a and at b, which, by definition and assumption, is greater than $\frac{\pi}{3}$.

Thus each long conjugate pair involves two vertices that belong to different portions of ∂P. There is only a constant number of pairs of such portions, and we now show that for each such pair there exist at most $O(n \log n)$ conjugate pairs 'connecting' between these two portions.

Let P_1, P_2 be the convex hulls of these respective portions, with P_1 regarded as lying counterclockwise to P_2. We analyze the number of conjugate pairs (a, b) of vertices, where a is a vertex of P_1 and b is a vertex of P_2.

Enumerate the vertices of P_1 in clockwise order as $p_1, \ldots, p_{n_1}$, and enumerate the vertices of P_2 in clockwise order as $q_1, \ldots, q_{n_2}$. Define an $n_1 \times n_2$ 0-1 matrix A such that $A_{ij} = 1$ iff the pair (p_i, q_j) are conjugate in P. We shall show that A does not contain the submatrix

$$(1) \qquad\qquad \begin{pmatrix} 1 & \cdot & 1 \\ \cdot & 1 & 1 \end{pmatrix} ,$$

involving any pair of rows and any triple of columns, in this order, where $\cdot$ denotes either 0 or 1. Hence, by the analysis of Füredi's [15] and of Bienstock and Györi [5], the number of ones in A is $O(n \log n)$. This, combined with the analysis given above, imply the assertion of the theorem.

Suppose to the contrary that A does contain such a submatrix, so there are two vertices a, b of P_1 and three vertices u, v, w of P_2, so that their clockwise order along ∂P is a, b, u, v, w, and the pairs (a, u), (a, w), (b, v), (b, w) are all conjugate in P; see Figure 6.

Figure 6: The configuration arising in the forbidden submatrix analysis

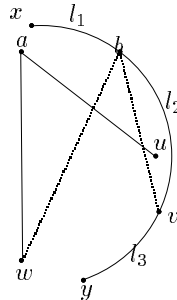Let Γ_0 be a copy of Γ that passes through b and v and avoids the interior of P. Let the counterclockwise and clockwise endpoints of Γ_0 be x and y, repsectively, and let $l_1 = \sigma(x, b)$,

$l_2 = \sigma(b, v)$, $l_3 = \sigma(v, y)$. We first claim that

$$
\begin{aligned}
l_1 &< \sigma(a, b) \,, \\
l_3 &< \sigma(v, w) \,.
\end{aligned}
$$

The second inequality follows directly from Corollary 2.2 (interchanging the role of clockwise and counterclockwise directions), since (b, w) and (b, v) are both conjugate pairs. The first inequality follows by applying Lemma 2.1 to the triple (a, b, u) of vertices of P. For this we only need to show that the ray from a to u intersects Γ_0 from the interior side of its disk D_0. This however is obvious, because, by convexity, the segment au and bv intersect, so the ray from a to u enters (through bv) the portion of D_0 bounded by the chord bv and by the middle arc Γ_2 of Γ_0, so it must exit this region by intersecting Γ_2 from inside D_0, as required. Hence $l_1 = \sigma(x, b) < \sigma(a, b)$.

We thus have

$$ l_1 + l_2 + l_3 < \sigma(a, b) + \sigma(b, v) + \sigma(v, w) \,. $$

Moreover, by Lemma 2.3,

$$ \sigma(b, v) + \sigma(v, w) < \sigma(b, w) \,, $$

and

$$ \sigma(a, b) + \sigma(b, w) < \sigma(a, w) \,, $$

since (b, w) and (a, w) are both conjugate pairs. Thus

$$ |\Gamma_0| = l_1 + l_2 + l_3 < \sigma(a, w) \le |\Gamma_0| \,. $$

This contradiction shows that the matrix A does not contain a submatrix of the form (1), and thus completes the proof of the theorem. $\square$

Remarks: (1) A close inspection of the preceding analysis shows that we can replace Γ_0 by any fixed arc smaller than a semicircle, and still get the bound $O(n \log n)$ on the number of conjugate pairs of vertices of a convex n-gon, with respect to such a smaller arc.

(2) If the vertices of P are not in general position, the number of conjugate pairs could be as large as $\Theta(n^2)$. This is the case, e.g., if all vertices lies on a circle of radius d.

2.2 Critical placements within a hippodrome

We now return to the original problem of bounding the number of critical placements of P in $H(e, d)$. Let p_i, p_j, p_k be the three vertices of P touching $\partial H(e, d)$ at a given critical placement. Clearly, one of the following cases must arise:

(i) Two of these vertices, say p_i and p_j, lie on one semicircle of $\partial H(e, d)$.

(ii) Two of these vertices, say p_i and p_j, lie on straight edges of $\partial H(e, d)$.

(iii) One vertex lies on a straight edge and one vertex lies on each semicircle of $\partial H(e, d)$.

In case (i) the two vertices p_i and p_j form a conjugate pair, by definition. Thus there are only $O(n \log n)$ such pairs. Moreover, we claim that each conjugate pair can participate in at most two critical placements of P. Indeed, suppose, with no loss of generailty, that p_i and p_j lie on the left semicircle of $H(e, d)$. Note that the locus of all placements of e at which both p_i and p_j lie on that semicircle is obtained by fixing the left endpoint a of e and by rotating e around that point through some angular interval I_0. Suppose that we keep e fixed and instead rotate the set P rigidly about a. Then each vertex $p_\ell \in P$ traces a circular arc which either is fully contained in the fixed hippodrome $H(e, d)$ or intersects its boundary in at most two points, as is easily verified. Thus each vertex p_ℓ defines an angular interval I_ℓ at which it lies inside $H(e, d)$ during this rotation, and the intersection

of I_0 with all these intervals, which is itself an angular interval, represents all placements of e, with its endpoint a held at a fixed point, at which the corresponding hippodrome contains P and the vertices p_i and p_j lie on the left semicircle of the hippodrome, and it is clear that only the two endpoints of this angular interval can represent critical placements of P of the type we seek here. Hence the number of critical placements of type (i) is $O(n \log n)$.

Critical placements of type (ii) are even easier to analyze. Suppose first that p_i and p_j lie on the same straight edge of $\partial H(e, d)$. Then it is clear that $p_i p_j$ must be an edge of P, so there are at most n such pairs. Moreover, each such pair induces at most two critical placements of P of the type considered here; these are obtained by translating P as far as possible to the left and to the right, while maintaining the double contact of p_i and p_j with the straight edge of $H(e, d)$, until a third point of P reaches $\partial H(e, d)$. A similar argument applies if p_i and p_j lie on different straight edges of $\partial H(e, d)$, observing that the number of such pairs is also $O(n)$, as follows easily from the standard 'rotating calipers' argument for a convex polygon. Hence the number of critical placements of type (ii) is only $O(n)$.

Finally, consider critical placements of e of type (iii). Assume, with no loss of generality, that at such a placement Z, p_i lies on the bottom straight edge of $H(e, d)$, p_j lies on the left semicircle and p_k lies on the right semicircle. Define

$$H^+(e, d) = \lambda^+(e) \oplus B_d \ ; \quad H^-(e, d) = \lambda^-(e) \oplus B_d \ ,$$

where $\lambda^+(e)$ (resp. $\lambda^-(e)$) is the ray that contains e and whose endpoint is the left (resp. right) endpoint of e. Note that, at placement Z, p_i and p_j lie on $\partial H^+(e, d)$ and all other points of P lie in the interior of that region (otherwise Z would also be a critical placement of one of the first two types). We now slide P continuously from Z, maintaining the contacts of p_i with the bottom edge and of p_j with the left semicircle of $H^+(e, d)$, and moving p_j counterclockwise, until either p_j reaches the counterclockwise endpoint of the semicircle, or a third point p_ℓ reaches $\partial H^+(e, d)$. In the former case $p_i p_j$ is an edge of P, so there are at most n such pairs. In the latter case, we have reached a critical placement of P inside $H^+(e, d)$, which of course must be of either type (i) or type (ii) (with respect to $H^+(e, d)$). Similar arguments to those used in the analysis of type (i) placements imply that the number of such critical placements is $O(n \log n)$. A fully symmetric argument shows that the number of pairs (p_i, p_k) participating in critical placements of type (iii) (with p_k lying on the right semicircle of the hippodrome) is also $O(n \log n)$.

Now fix a vertex p_i of P, constrain p_i to lie on, say the bottom edge of the hippodrome, and consider the set Q_i^+ (resp. Q_i^-) of all vertices p_j (resp. p_k) of P that participate with p_i in a critical placement of type (iii) and lie on the left (resp. right) semicircle of $H(e, d)$ at such a placement, as defined above. The preceding analysis implies that

$$\sum_i (|Q_i^+| + |Q_i^-|) = O(n \log n) \ .$$

For each $p_j \in Q_i^+$ define a function $f_j(\theta)$, as follows. For each orientation θ of e, at which e is parallel to a line supporting P at p_i, consider a placement of $H(e, d)$ at which e is at orientation θ, p_i is on the bottom edge and p_j touches the left semicircle of $H(e, d)$ (see Figure 7). We define $f_j(\theta)$ to be the distance from p_i to the left endpoint of the bottom edge of $H(e, d)$. If no such placement exists we put $f_j(\theta) = -\infty$. For each $p_k \in Q_i^-$ we define a function $g_k(\theta)$ in a fully symmetric manner, still measuring the distance from p_i to the left endpoint of the bottom edge of $H(e, d)$, but requiring p_k to lie on the right semicircle of $H(e, d)$; again we put $g_k(\theta) = +\infty$ if no corresponding placement of e exists.

The following is an easy consequence of the definition of these functions: Suppose that $Z = (t, \theta)$ is a free placement of P inside $H(e, d)$ at which p_i lies on the bottom edge of the hippodrome, where θ is the orientation of e and t is the distance from p_i to the left endpoint of the bottom edge. Then we must have

$$\max_{p_j \in Q_i^+} f_j(\theta) \le t \le \min_{p_k \in Q_i^-} g_k(\theta) \ .$$
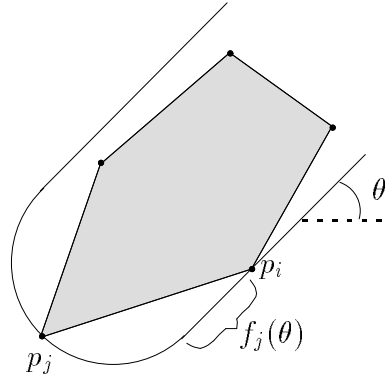
Figure 7: The definition of $f_j(\theta)$

Moreover, a triple (p_i, p_j, p_k) of vertices of P induces a critical placement of type (iii) of the form considered above if and only if the parameters (t, θ) of that placement represent a point of intersection between the two envelopes

$$\Phi_i(\theta) = \max_{p_j \in Q_i^+} f_j(\theta), \quad \Psi_i(\theta) = \min_{p_k \in Q_i^-} g_k(\theta).$$

Using standard arguments, it is easily seen that the number of such intersection points is proportional to the number of vertices of the envelopes Φ_i, Ψ_i. However, any such vertex is easily seen to represent a critical placement of P in either $H^+(e, d)$ or $H^-(e, d)$, and we have already argued that the number of such placements is $O(n \log n)$ (over all possible choices of p_i). We thus conclude that the number of critical placements of type (iii), and thus also the overall number of critical placements of P, is $O(n \log n)$:

Theorem 2.5 *Under the general position assumption, the number of critical free placements of a convex n-gon inside a given hippodrome is $O(n \log n)$.*

Remark. With some care, Theorem 2.5 also holds without the general position assumption. For this we need to count only critical placements in which the vertices of the polygon that lie on $\partial H(e, d)$ are not all contained in a single semicircle.

3   Computing All Critical Placements

In this section we exploit the improved combinatorial bounds on the number of critical placements to design an efficient algorithm for computing all these placements. (Note however that in our application it suffices to only determine whether there exists at least one such critical placement.) Let S, P, e, and d be as above. Our goal is to find all placements Z of e at which the hippodrome $H(e(Z), d)$ contains P and three vertices of P lie on its boundary; here $e(Z)$ is the set of points occupied by e at placement Z.

Before plunging into our analysis, here is a brief explanation of the structures that the algorithm attempts to construct. The space Π of all placements Z of e is 3-dimensional. For each vertex p_i of P, we can draw in Π the surface Σ_i which is the locus of all placements Z at which p_i lies on $\partial H(e(Z), d)$. The surface Σ_i partitions Π into the regions Σ_i^+, Σ_i^-, of all placements Z where p_i lies, respectively, in the interior or in the exterior of $H(e(Z), d)$. Our goal is to compute $\Pi^0 = \cap_{i=1}^n \Sigma_i^+$ (or, rather, to determine whether this intersection is nonempty). This intersection is the union of

certain cells in the 3-dimensional arrangement of the surfaces Σ_i. The problem is difficult to solve efficiently in this abstract setting, so we attack it in a more careful manner. We first show that, with an appropriate parametrization, Π^0 can be represented as the region lying between the lower envelope of one collection of (partial) bivariate functions and the upper envelope of another such collection. Each of the envelopes has only near-linear complexity, and can be constructed in near-linear time by a standard, ε-net-based, technique. The effieicnt construction of the region Π^0 itself is more complicated, and requires the use of additional machinery, as described below.

In more detail, the algorithm proceeds as follows. For each vertex p_i of P we define two partially defined bivariate functions, F_i and G_i. Both functions are defined in the dual plane, where each point (ξ_1, ξ_2) is the dual of the line $\lambda(\xi_1, \xi_2) : y = \xi_1 x + \xi_2$. For each point (ξ_1, ξ_2) and each real number t, let ρ_t^+ (resp. ρ_t^-) denote the right half (resp. left half) of $\lambda(\xi_1, \xi_2)$, delimited at the point $(t, \xi_1 t + \xi_2)$. We define $F_i(\xi_1, \xi_2)$ (resp. $G_i(\xi_1, \xi_2)$) to be the value of t at which the boundary of the 'half-hippodrome' $H(\rho_t^+, d) = \rho_t^+ \oplus B_d$ (resp. $H(\rho_t^-, d) = \rho_t^- \oplus B_d$) passes through p_i; see Figure 8. If p_i lies at distance d from $\lambda(\xi_1, \xi_2)$ then the functions F_i, G_i are not well defined; the locus of (ξ_1, ξ_2) where this occurs is the boundary of the common domain of F_i and of G_i. If no such t exists (that is, p_i lies at distance $> d$ from the line $\lambda(\xi_1, \xi_2)$), we put $F_i(\xi_1, \xi_2) = -\infty$, $G_i(\xi_1, \xi_2) = +\infty$.
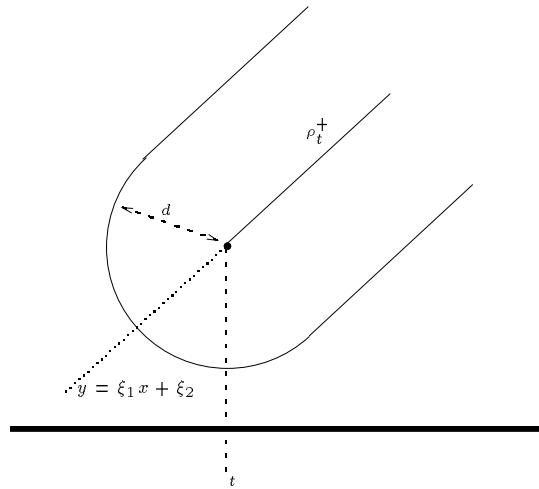

Figure 8: The rightward-directed semi-hippodrome $H(\rho_t^+, d)$

Let us parametrize each placement Z of e by the triple (ξ_1, ξ_2, t), where (ξ_1, ξ_2) is the point dual to the line containing $e(Z)$, and where t is the x-coordinate of the right endpoint b of e at placement Z. Note that the x-coordinate t' of the left endpoint of e at this placement satisfies $t' = t - 1/\sqrt{1 + \xi_1^2}$. (This parametrization excludes placements Z for which $e(Z)$ is vertical; these placements can be handled by a different, much simpler procedure, whose description is omitted here.)

It is now immediate from the definitions of the functions F_i, G_i that $Z = (\xi_1, \xi_2, t)$ is a *valid* placement of e (i.e., a placement at which $H(e(Z), d)$ fully contains P) if and only if

$$(1) \qquad \max_{1 \le i \le n} G_i(\xi_1, \xi_2) \le t \le \min_{1 \le i \le n} F_i(\xi_1, \xi_2) + \frac{1}{\sqrt{1 + \xi_1^2}} \ .$$

Let us denote by V the set of all valid placements of e. Note that a placement $Z \in V$ is critical if and only if Z is a *vertex* of V, that is, Z lies on the boundary of V and either (a) at least three inequalities in (1) (or exactly three, if we assume general position) become equalities, or (b) for some $j = 1, 2$, or 3, j points of P lie at distance d from $\lambda(\xi_1, \xi_2)$, and $3 - j$ other points of P induce an equality in (1). Thus our goal is to compute all the vertices of V.

There are some technical difficulties in computing the region V directly, so we use the following roundabout way. We first compute the lower envelope $F(\xi_1, \xi_2) = \min_i F_i(\xi_1, \xi_2)$, using a standard ε-net approach, similar to those described in [24, 26]. That is, we fix some sufficiently large but constant integer parameter r, and construct a $(1/r)$-net $N \subseteq S$, of size $O(r \log r)$, for a range space $(S, \mathcal{R})$, defined in terms of a set $\mathcal{R}$ of ranges that will be specified shortly. We construct the lower envelope

$$F^{(N)}(\xi_1, \xi_2) = \min \{F_i(\xi_1, \xi_2) \mid p_i \in N\} ,$$

using any brute-force technique. The projection of $F^{(N)}$ onto the $\xi_1\xi_2$-plane is a planar map M whose combinatorial complexity, by Theorem 2.4, is only $O(r \log^2 r)$. We construct the *vertical decomposition* of M, by drawing a ξ_2-vertical segment from each vertex of M, and from each point with ξ_2 vertical tangency on any edge of M, and extend each such segment until it hits another edge of M; see, e.g., [26] for more details. For each of the $O(r \log^2 r)$ resulting 'pseudo-trapezoidal' cells c, we define

$$(2) \qquad c^* = \{(\xi_1, \xi_2, t) \mid (\xi_1, \xi_2) \in c \ \text{and} \ t \leq F^{(N)}(\xi_1, \xi_2)\} .$$

Since $F^{(N)}$ is attained over c by a single function F_i, for some $p_i \in N$, the region c^* has constant description complexity, and is defined by a constant number of points of N. The set $\mathcal{R}$ of ranges, with respect to which N has to be a $(1/r)$-net, is defined to consist of all subsets $P(c^*)$ of points of P, where c^* is a spatial region of the form defined in (2), and $P(c^*)$ is the set of all points p_i of P such that the graphs of their associated functions F_i cross c^*. (We omit here more precise details concerning the structure of the regions c^*, which can be easily worked out from the definition in (2). It is easily checked that the resulting range space has finite VC-dimension, which implies the existence of $(1/r)$-nets of size $O(r \log r)$, and, using the technique of [20], such nets can be constructed in $O(n)$ time (for constant r).

Note that, by construction, none of the cells c^*, defined above for $F^{(N)}$, is crossed by the graph of any function F_i, for $p_i \in N$. It thus follows that $|P(c^*)| \leq n/r$ for all cells c of M. The sets $P(c^*)$ can be constructed in $O(n)$ time by a brute-force technique, since there is a constant number of regions c^* and each of them has constant complexity. Hence we can compute recursively the lower envelope F over each cell c of M separately, and then 'glue' together the resulting portions of F. It follows that the maximum time complexity $T(n)$, required by this procedure for an input set of n points, satisfies the recurrence:

$$T(n) \leq cr \log^2 r T(n/r) + O(n) ,$$

where c is some absolute constant, and the constant of proportionality in the term $O(n)$ depends on r. The solution of this recurrence is easily seen to be $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε (and increases to ∞ as ε tends to 0).

Using a completely symmetric procedure, we can also compute the upper envelope $G(\xi_1, \xi_2) = \max_i G_i(\xi_1, \xi_2)$, in time $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$. We also modify $F(\xi_1, \xi_2)$ by adding to it the term $1/\sqrt{1 + \xi_1^2}$, in accordance with (1).

Remark: We could have applied the same ε-net approach to the explicit calculation of the region V. However, as seen from the preceding analysis, the efficiency of the resulting algorithm depends on the number of subcells of constant description complexity in an appropriate decomposition of the region $V^{(N)}$, defined in complete analogy to V but only in terms of the functions F_i, G_i corresponding to the points p_i in the net N. So far we have failed to obtain such a decomposition with only a near-linear number of cells. One possibility is to apply to $V^{(N)}$ the standard 3-dimensional *vertical decomposition* (as in [8]). The number of its subcells is easily seen to be proportional to the number of quadruples (a, b, c, d) of vertices of P, for which there exists a hippodrome $H(e', d)$, where e' is some segment of length > 1, such that $P \subseteq H(e', d)$, a and b lie on one semicircle of that hippodrome, and c and d lie on the other semicircle. A near-linear bound on the number of such quadruples would lead to an efficient algorithm for computing V directly, in which case the remainder of the analysis in this section becomes unnecessary. See also [1] for an analysis of the complexity of regions like V in a more general setting, where only a near-quadratic bound is obtained.

Next we need to 'interact' the two envelopes in order to find the vertices of V. We first observe that vertices of V that are also vertices of either F or G are easy to find: After computing the full envelopes F, G, we process their $\xi_1\xi_2$-projections for efficient point location (e.g., as in [11]), and locate the $\xi_1\xi_2$-projections of the vertices of F in the map of G, and vice versa. This allows us to test, in additional $O(1)$ time per vertex (using (1)), whether that vertex lies in the 'correct' side of the other envelope, and is thus a vertex of V. The cost of this step is easily seen to be $O(n \log^2 n)$.

It is more difficult to compute the vertices of V at which the envelopes F and G meet. Let Z be the placement of e corresponding to such a vertex, and let p_i, p_j, p_k be the three vertices of P lying on $\partial H(e(Z), d)$. Then, as is easily seen, Z must be of one of the following types:

(a) At least one of the points p_i, p_j, p_k lies on a straight edge of $\partial H(e(Z), d)$.

(b) Two of these points lie on one semicircle bounding $H(e(Z), d)$ and the third point lies on the other semicircle.

Consider case (a) first. If two (or three) of the points p_i, p_j, p_k lie on straight edges of $\partial H(e(Z), d)$, then, as is easily seen, Z appears as a vertex of at least one of the envelopes F, G (lying over the boundaries of the domains of two of the functions), so the point-location technique described above will have already identified such placements. Otherwise, it must be the case that one point, say p_i, lies on a straight edge s of $\partial H(e(Z), d)$, another point p_j lies on the right semicircle of $H(e(Z), d)$ and the third point p_k lies on the left semicircle. It is easily verified that the pair (p_i, p_j) defines an edge of the envelope F (it is the locus of points (ξ_1, ξ_2, t) at which p_i lies at distance d from $\lambda(\xi_1, \xi_2)$ and p_j lies on the left semicircle of $H(\rho_t^+, d)$; the $\xi_1\xi_2$-projection of that edge is a portion of the boundary of the domain of F_i), and the pair (p_i, p_k) defines a similar edge of the envelope G (projecting onto a similar boundary portion). Let $E(p_i)$ denote the set of all points $p \in P$ such that the pair (p_i, p) defines an edge of either F or G, representing placements Z at which p_i lies on a straight edge of $\partial H(e(Z), d)$. Note that $\sum_i |E(p_i)| = O(n \log n)$, and that all the sets $E(p_i)$ can be constructed in $O(n \log n)$ time, by simply scanning all relevant edges of the envelopes F and G. Now let $\gamma(p_i)$ denote the curve in the $\xi_1\xi_2$-plane, representing the locus of all lines $\lambda(\xi_1, \xi_2)$ that are tangent to the circle of radius d about p_i (the curve $\gamma(p_i)$ contains the boundary arc of the domain of F_i over which Z lies). Define F' (resp. G') to be the envelope F (resp. G) restricted over $\gamma(p_i)$. Note that the complexity of F' and of G' is $O(|E(p_i)|)$, since the vertices of these envelopes correspond to conjugate pairs at which p_i participates while touching s, and each such pair appears at most once.

It is easily checked that the placements Z that we seek appear as points of intersection between the graphs of F' and of G'. Since these functions are univariate, these intersections can all be found in time proportional to the total complexity of F' and of G', namely in $O(|E(p_i)|)$ time. Thus all placements Z in the present subcase can be computed in time $O(n \log^2 n)$. (The analysis given here is very similar to that given at the end of Section 2.2 for critical placements of type (iii).)

Next consider case (b). Suppose, with no loss of generailty, that p_i and p_j lie on the left semicircle of $H(e(Z), d)$ and that p_k lies on the right semicircle. Note that the pair (p_i, p_j) defines an edge δ of F, so that moving along δ corresponds to rotating e about its left endpoint a. We now argue in a manner similar to the analysis of type (i) critical placements given in Section 2.2. Suppose that we keep e fixed and instead rotate the set P rigidly about a. Then each point $p_i \in P$ traces a circular arc which either is fully contained in the fixed hippodrome $H(e, d)$ or intersects its boundary in at most two points. Thus each point p_i defines an angular interval I_i at which it lies inside $H(e, d)$ during this rotation, and the intersection of these intervals, which is itself an angular interval, represents all placements Z of e, with its endpoint a held at a fixed point, at which $H(e(Z), d)$ contains P; the two endpoints of this angular interval represent vertices of V of the type we seek here. Note that this argument implies that there are at most two vertices of this type along the arc δ.

Since we already have all these edges δ of F at our disposal, our strategy is to process the vertices of P into some data structure that we will later query with each of the curves δ to find the vertices of the type we seek that lie along δ. For this we use another parametrization of the space Π of

placements of e, representing each placement by the triple (x, y, θ), where (x, y) is the position of the left endpoint a of e and $-\frac{\pi}{2} \le \theta \le \frac{\pi}{2}$ is its orientation. For each point $p_i \in P$ we define the following two partial bivariate functions:

- $\phi_i(x, y) =$ the smallest of the two orientations θ at which p_i lies on $\partial H(e((x, y, \theta)), d)$;

- $\psi_i(x, y) =$ the largest of these two orientations;

if no such orientations exist at (x, y) then either p_i lies inside the hippodrome for all possible θ, or p_i always lies outside that hippodrome. In the former case we define $\phi_i(x, y) = -\frac{\pi}{2}$, $\psi_i(x, y) = \frac{\pi}{2}$, and in the latter case we define $\phi_i(x, y) = \frac{\pi}{2}$, $\psi_i(x, y) = -\frac{\pi}{2}$.

Note that, if a is held fixed at the point (x, y), the angular interval I_i defined above is $\phi_i(x, y) \le \theta \le \psi_i(x, y)$. This suggests the following approach. We compute the upper envelope $\Phi(x, y) = \max_i \phi_i(x, y)$ and the lower envelope $\Psi(x, y) = \min_i \psi_i(x, y)$. Note that a vertex of $\Psi(x, y)$ represents a critical placement of P in a 'half-hippodrome', defined as the region $H^* = e^* \oplus B_d$, where $e^* = \{(x, y) \mid 0 \le x \le 1, \ y \ge 0\}$, so that three vertices of P lie on the union of the bottom edge and the right quarter-circle of H^*; see Figure 9 for an illustration. A symmetric property holds for the vertices of $\Phi(x, y)$.
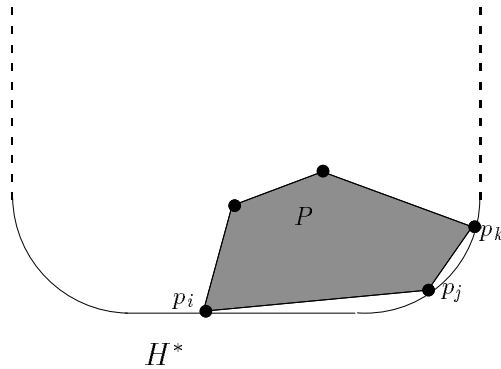

Figure 9: The half-hippodrome used in the construction of Ψ

Now at any such critical placement, either two vertices of P lie on the bottom edge, or two vertices lie on the right quarter-circle of H^*. In the former case, arguing as above, we conclude that there are at most n such placements. In the latter case, we can apply the analysis of Section 2, replacing the semicircle Γ_0 used there by a quarter-circle of the same radius. As observed at the end of Section 2.1, the analysis given there continues to apply to such smaller arcs, and thus the number of placements under consideration is $O(n \log n)$.

Thus the complexity of the envelopes Φ, Ψ, is $O(n \log n)$, so we can apply the same procedure, described earlier in this section, to compute these envelopes in time $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$. We project the envelopes onto the $\xi_1 \xi_2$-plane, and process these projections for efficient point location. Now each edge δ of F of the type discussed above defines a point (x_δ, y_δ), which is the common location of the left endpoint of e for all placements in δ. We thus locate (x_δ, y_δ) in the projections of Φ and Ψ, and consequently obtain $\Phi(x_\delta, y_\delta)$ and $\Psi(x_\delta, y_\delta)$. If $\Phi(x_\delta, y_\delta) \le \Psi(x_\delta, y_\delta)$, then $(x_\delta, y_\delta, \Phi(x_\delta, y_\delta))$ and $(x_\delta, y_\delta, \Psi(x_\delta, y_\delta))$ are the two vertices of V along δ. If $\Phi(x_\delta, y_\delta) > \Psi(x_\delta, y_\delta)$, then there are no vertices of V along δ (and δ is fully disjoint from ∂V). We repeat this process for all edges δ of F, and then apply a symmetric calculation to the edges of G, using a parametrization that depends on the location of the right endpoint, rather than the left endpoint, of e. The total cost of all these point locations is $O(n \log^2 n)$, which is dominated by the cost of constructing the envelopes Φ, Ψ, and their symmetric counterparts. We thus conclude:

Theorem 3.1   *(a) All critical placements of a convex n-gon in a given hippodrome can be computed in time $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε.*

(b) Consequently, given a set S of n points in the plane, a segment e and a distance $d > 0$, one can determine, in time $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, whether there exists a placement of e at which its maximum distance from any point of S is $\leq d$.

4   Computing the Segment Center

We now apply the parametric searching paradigm to the algorithm given in Section 3, to compute the desired center location of e. For this we first need to design a parallel version of this algorithm, in Valiant's comparison model [28]. See, e.g., [3] for more details concerning parametric searching.

Fortunately, the preceding algorithm is very easy to parallelize. Consider for example the construction of the lower envelope $F(\xi_1, \xi_2)$. This proceeds in $O(\log n)$ stages, where in each stage we further reduce the problem size by a factor of r, but increase the number of subproblems by a factor of $O(r \log^2 r)$. However, all these subproblems can be solved in parallel. At each stage we need to compute a $(1/r)$-net for an approrpriate set of points. This is done using the parallel version of the technique of [20], as presented in [9]. This version takes $O(\log n)$ parallel steps, using a linear number of processors. We thus obtain a parallel procedure that runs in $O(\log^2 n)$ parallel stages and uses $O(n^{1+\varepsilon})$ processors, for any $\varepsilon > 0$, as is easily checked.

The next stage is to process such an envelope for fast point location. Rather than applying known (but complicated) parallel point location schemes (such as that in [27]), we observe that the output of the $O(\log n)$ stages of the construction of F yields a data structure that already supports point locations: given a query point Z, locate it (in constant time) in the vertical decomposition of the projection of the 'sample envelope' $F^{(N)}$, and continue the search at the recursive substructure computed for the cell that contains Z. Thus point location can be done in $O(\log n)$ time, at no extra preprocessing cost, and we perform all these point locations in parallel.

Similar techniques can be applied to all other stages of the algorithm. We leave it to the reader to verify that this results in a parallel version of the algorithm that takes $O(\log^2 n)$ time and uses $O(n^{1+\varepsilon})$ processors.

Plugging this parallel algorithm into the parametric searching paradigm, we readily obtain the desired algorithm for computing the segment center of the given set S, whose running time is as asserted in Theorem 1.1.

5   Conclusion and Open Problems

In this paper we have drastically recduced the complexity of computing the segment center of a given set of n points in the plane, from near-quadratic to near-linear in n. The highlight of our analysis is a near-linear bound on the number of critical placements of a convex polygon inside a hippodrome, combined with several nontrivial computational tricks. There are evidences that better running time might be obtained; However, the predicted improvement does not seem to justify the effort.

There are still some interesting open problems to consider. The main one is whether our bound on the number of conjugate pairs of vertices of a convex n-gon is tight, and, if not, to tighten it further. Another open problem is to further improve the running time of the algorithm. For instance, can the problem be solved more efficiently without the use of parametric searching? Can the running time be improved to $O(n \operatorname{polylog}(n))$ time?

Other natural extensions of the problem include:

- Given a set S of n points in the plane, a segment e, and a parameter $k \leq n$, find the smallest value of d for which there exists a copy e' of e and a subset $S' \subseteq S$ of cardinality k, such that $S' \subseteq H(e', d)$. This problem is a natural extension of the well-studied problem of computing the smallest disc enclosing k points of S; see [10, 12, 14, 21]. A recent attack on this problem is given in [2], (See also [13]).

- The segment center problem is actually a special case of the more general problem of computing the *one-directional Hausdorff distance*, under euclidean motion, between two sets of objects, which can be stated as follows (see [16]):

  Given two sets, S_1, S_2, of objects in the plane, find the smallest value of d such that the objects of S_1 can be translated and rotated (rigidly) into a placement where every object of S_2 is within distance d from some object of S_1.

  Hence an obvious open problem is whether our technique can also be applied to solve problems of this kind efficiently.

Acknowledgment.

References

[1] The overlay of lower envelopes in 3-space and its applications, *Proc. 11th ACM Symp. on Computational Geometry* (1995), 182–189. (Also to appear in *Discrete Comput. Geom.*)

[2] P. Agarwal, A. Efrat and M. Sharir, Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications, *Proc. 11th ACM Symp. on Computational Geometry* (1995), 39–50.

[3] P.K. Agarwal, A. Efrat, M. Sharir and S. Toledo, Computing a segment-center for a planar point set, *J. Algorithms* 15 (1993), 314–323.

[4] B. Bhattacharya, J. Czyzowicz, P. Egyed, I. Stojmenovic, G. Toussaint and J. Urrutia, Computing shortest transversals of sets, *Proc. 7th ACM Symp. on Computational Geometry*, 1991, 71–80.

[5] D. Bienstock and E. Györi, An extremal problem on sparse 0-1 matrices, *SIAM J. Discrete Math.* 4 (1991), 17–27.

[6] V. Capoyleas and J. Pach, A Turán-type problem on chords of a convex polygon, *J. Combin. Theory, Ser. B* 56 (1992), 9–15.

[7] B. Chazelle, The polygon containment problem, in *Advances in Computing Research, Vol. I: Computational Geometry,* (F.P. Preparata, Ed.), JAI Press, Greenwich, Connecticut (1983), pp. 1–33.

[8] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, A singly exponential stratification scheme for real semi–algebraic varieties and its applications, *Theoretical Computer Science* 84 (1991), 77–105.

[9] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, Diameter, width, closest line pair, and parametric searching, *Discrete Comput. Geom.* 10 (1993), 183–196.

[10] A. Datta, H.-P. Lenhof, C. Schwarz, and M. Smid. Static and dynamic algorithms for k-point clustering problems. In *Proc. 3rd Workshop Algorithms Data Struct.*, Lecture Notes in Computer Science, 1993, 265–276..

[11] H. Edelsbrunner, L. Guibas and J. Stolfi, Optimal point location in a monotone subdivision, *SIAM J. Comput.* 15 (1986), 317–340.

[12] A. Efrat, M. Sharir and A. Ziv, Computing the smallest k-enclosing circle and related problems, *Comput. Geom. Theory Appls.* 4 (1994), 119–136.

[13] A. Efrat, *Location Problems in Geometric Optimization*, Ph.D. Dissertation, Tel Aviv University, in preparation.

[14] D. Eppstein and J. Erickson. New algorithms for minimum measure simplices and one-dimensional weighted Voronoi diagrams. Tech. Report 92-55, Dept. Inform. Comput. Sci., Univ. California, Irvine, CA, 1992.

[15] Z. Füredi, The maximum number of unit distances in a convex n-gon, *J. Combin. Theory, Ser. A* 55 (1990), 316–320.

[16] L.P. Chew, M. Goodrich, D. Huttenlocher, K. Kedem, J. Kleinberg and D. Kravetz, Geometric pattern matching under euclidean motion, *Proc. 5th Canadian Conf. on Computatoonal Geometry*, 1993, 151–156.

[17] H. Imai, D. T. Lee, and C. Yang, 1-Segment center covering problems, *ORSA J. Computing* 4 (1992), 426–434.

[18] D. T. Lee and Y. Wu, Geometric complexity of some location problems, *Algorithmica* 1 (1986), 193–211.

[19] D. Leven and M. Sharir, On the number of critical free contacts of a convex polygonal object moving in two-dimensional polygonal space, *Discrete Comput. Geom.* 2 (1987), 255–270.

[20] J. Matoušek, Approximations and optimal geometric divide-and-conquer, *Proc. 23rd ACM Symp. on Theory of Computing* (1991), 506–511.

[21] J. Matoušek, On enclosing k points by a circle, manuscript, 1993.

[22] J. Matoušek and O. Schwarzkopf, A deterministic algorithm for the three-dimensional diameter problem, *Proc. 25th Annu. ACM Sympos. Theory of Computing* (1993), 478–484.

[23] N. Megiddo, Applying parallel computation algorithms in the design of serial algorithms, *J. ACM* 30 (1983), 852–865.

[24] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.

[25] M. Sharir and S. Toledo, Extremal polygon containment problems, *Comp. Geom. Theory Appls.* 4 (1994), 99–118.

[26] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, to appear.

[27] R. Tamassia and J. Vitter, Parallel transitive closure and point location in planar structures, *SIAM J. Computing* 20 (1991), 708–725.

[28] L. Valiant, Parallelism in comparison problems, *SIAM J. Computing* 4 (1975), 348–355.