On the Union of κ-Curved Objects*

Alon Efrat † Matthew J. Katz ‡

July 16, 1998

Abstract

A (not necessarily convex) object C in the plane is *κ-curved* for some constant κ, $\kappa < 1$, if it has constant description complexity, and for each point p on the boundary of C, one can place a disk B whose boundary passes through p, its radius is $\kappa \cdot diam(C)$ and it is contained in C. We prove that the combinatorial complexity of the boundary of the union of a set $\mathcal{C}$ of n κ-curved objects (e.g., fat ellipses or rounded hearts) is $O(\lambda_s(n) \log n)$, for some constant s. We also describe an efficient dynamic data structure for point location queries for $\mathcal{C}$.

1 Introduction

Let C be a (not necessarily convex) object in the plane, and let κ, $\kappa < 1$, be a constant. We say that C is *κ-curved* if

(i) C has constant description complexity s_0. (This implies, for example, that there are at most s_0 local minima or maxima on its boundary, and that the number of intersection points between the boundary of C and the boundary of another object C' with constant description complexity s_0 is at most s_0 (assuming general position).)

*Work by the second author has been supported by the Israel Science Foundation founded by The Israel Academy of Sciences and Humanities.

†Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. *alone@math.tau.ac.il, http://www.math.tau.ac.il/~alone*

‡Department of Mathematics and Computer Science, Ben-Gurion University of the Negev, Beer-Sheva 84105, Israel. *matya@cs.bgu.ac.il, http://www.cs.bgu.ac.il/~matya*

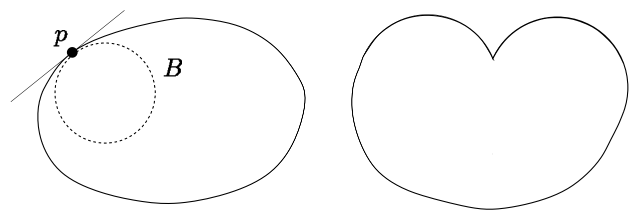

Figure 1: Two κ-curved objects

(ii) For each point p on the boundary of C, we can place a disk B whose boundary passes through p, its radius is $\kappa \cdot diam(C)$, and it is contained in C; see Figure 1. We say that the radius $\kappa \cdot diam(C)$ is the *critical radius* of C, and that the disk B is a *critical disk* of C at p.

The second condition is similar to bounding the curvature of the boundary of C, but is more general (see Figure 1). It can be illustrated as follows: Imagine a car moving along the boundary of C such that the interior of C is to its left. Then the car is allowed to make very sharp right turns, but when turning left, the radius of the turn is bounded from below by some fraction of the diameter of C.

Regarding the second condition, if C has a tangent at a point p on its boundary, then there exists only one critical disk B of C at p, and its tangent at p coincides with the tangent of C at p. However, C may have a constant number of points on its boundary at which the tangent is not defined. At these points, though, C does have a left and a right tangent.

In this paper we prove that for any $0 \leq \kappa \leq 1$ and for any set $\mathcal{C} = \{C_1, \ldots, C_n\}$ of n κ-curved objects, the number of vertices on the boundary of the union U of the objects in $\mathcal{C}$ is only $O(\lambda_s(n) \log n)$, for some constant s. Alon says: Uncomment or something ment

We say that an object C is α-fat, for a constant $\alpha > 1$, if $r_1/r_2 \leq \alpha$, where r_1 is the radius of a smallest disk containing C, and r_2 is the radius of a largest disk contained in C. Obviously, a κ-curved object is α-fat for an appropriate constant α, but the opposite statement is false. Fat objects received much attention in recent years. One of the first papers on fat objects is by Matoušek et al. [10] who showed that a set of n triangles determines only a linear number of "holes," and that the number of vertices on the boundary of its union is only $O(n \log \log n)$. Since then many authors considered various definitions of fatness (which are all more or less equivalent—at least for convex objects), and obtained either interesting combinatorial results or efficient geometric algorithms for various classes of fat objects (see e.g. [3, 6, 7, 8, 9, 12, 14, 15]). However, the question which properties suffice so that the number of vertices on the boundary of the union of a set of planar objects having these properties is always subquadratic remained open for quite a few years. Recently, Efrat and Sharir [5] showed that if the objects are convex, fat, and the boundaries of each pair intersect at most a constant number of times, then the boundary of their union consists of only $O(n^{1+\varepsilon})$ vertices, for any constant $\varepsilon > 0$ [1]. In a preliminary version of their paper, it was shown that if, in addition, the object have bounded curvature and are more or less of the same size, then the number of vertices on the union's boundary is only $O(\lambda_s(n))$, for some constant s.

Our result improves upon the result of [5] for convex κ-curved objects such as fat ellipses, and complements it for objects that are non-convex (but κ-curved). We prove the following theorem.

Theorem 1.1 *The combinatorial complexity of the boundary of the union of n κ-curved objects is $O(\lambda_s(n) \log n)$, for some constant s.*

In the proof we use a well known data structure, namely, a segment tree, and its properties. We project the input κ-curved objects on the y-axis, and construct a segment tree $\mathcal{T}$ for these projections. We then insert the objects into $\mathcal{T}$ according to their projection on the y-axis. As usual, we associate with a node μ of $\mathcal{T}$ its canonical y-interval, which we think of as an horizontal slab. Now, roughly speaking, the vertices on the boundary of the union of the input objects are distributed among the nodes of $\mathcal{T}$, so that, if a vertex w ends up at a node μ, then w lies in the canonical slab of μ and is formed by a pair of objects that are stored at μ. (The objects that are stored at μ consist of the objects in the canonical subset of μ and the objects in the canonical subsets of all the descendants of μ.) By proving a connection between the number of vertices that end up at μ and the number of objects that are stored at μ, and by summing over all nodes in $\mathcal{T}$, we obtain the claimed bound.

[1] Throughout the paper, ε stands for a positive constant which can be chosen arbitrarily small with an appropriate choice of other constants of the big-O notation.

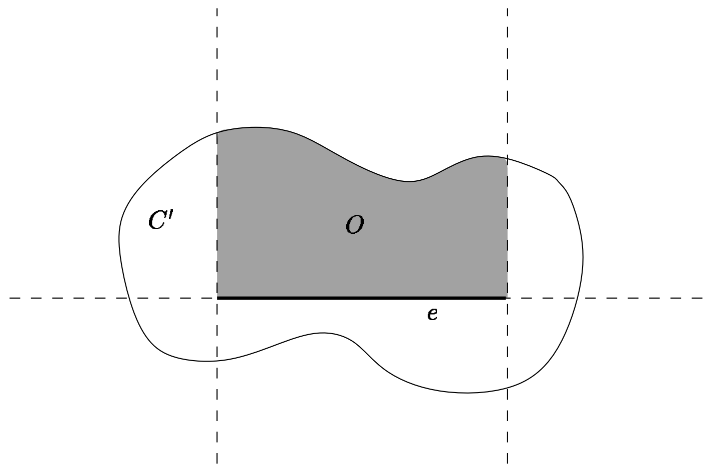

Figure 2: C' is function-defined with respect to e (from above)

The paper is organized as follows. In Section 2 we establish an auxiliary result, which is later used (Claim 3.2) in the proof of Theorem 1.1. The proof of Theorem 1.1 is given in Section 3, except for the proof of the, so called, key lemma, which is stated in this section and proven in Section 4.

2 Partitioning the boundary of an object

Let ρ be a horizontal strip of width δ that is divided into (axis-aligned) squares $\Sigma = \{\sigma_1, \sigma_2, \ldots\}$ of edge-length δ. Let ℓ_{top} and ℓ_{bottom} denote the top and bottom horizontal lines defining ρ. Let C be a κ-curved object whose diameter is at least $3\delta/\kappa$. (The radius of a critical disk of C is therefore at least 3δ). In this section we show that it is possible to obtain from C a constant number of smaller objects (called *parts*), such that (i) each part is contained in C and has some desirable properties, and (ii) for each point p on $\rho \cap \partial C$, there exists a part such that p lies on its boundary. We use this result in the proof of Theorem 1.1.

We need the following definition.

Definition: Consider an object C' in the plane, and a horizontal segment e. We cut off the part of C' that lies below the line containing e, and the

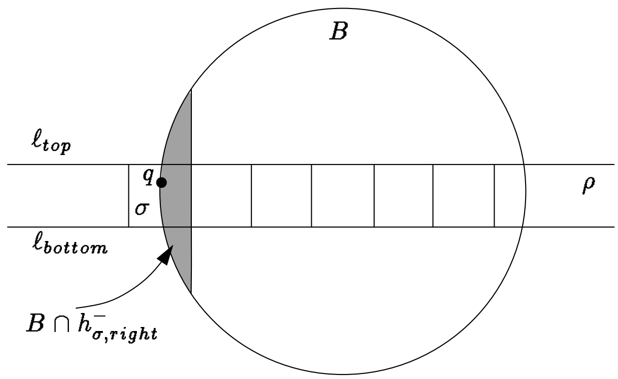

Figure 3: $B \cap h^-_{\sigma,right}$ is function defined with respect to $\ell_{\sigma,right}$

parts of C' that lie outside the vertical strip defined by e. Let O denote the remaining part of C'. We say that C' is *function-defined* with respect to e (from above), if for every point p on e, the intersection of O with the vertical ray emanating from e and directed upwards is a closed segment whose bottom endpoint is p (see Figure 2). In other words, the top boundary of O is the graph of some function defined on e, and O is the area enclosed between this graph and the segment e. We define in an analogous manner the statements: C' is *function-defined* with respect to e from below, or, for a vertical segment e, C' is *function-defined* with respect to e from the left (alternatively, from the right).

For each point p on ∂C, let $\mathcal{B}(p)$ denote the collection of disks of radius $\kappa \cdot diam(C)$ that are contained in C and whose boundaries pass through p. (As mentioned, if ∂C has a tangent at p, then $\mathcal{B}(p)$ consists of a single disk.) For a disk $B \in \mathcal{B}(p)$, let B' denote the disk of radius 3δ obtained from B by moving its center towards p while maintaining the contact with p. Let h^+_{bottom} (resp. h^-_{top}) denote the halfplane bounded by ℓ_{bottom} (resp. ℓ_{top}) and lying above ℓ_{bottom} (resp. below ℓ_{top}). Similarly, for a square $\sigma \in \Sigma$, let $\ell_{\sigma,left}$ and $\ell_{\sigma,right}$ denote the lines containing the left and right edges of σ, respectively, and let $h^+_{\sigma,left}$ (resp. $h^-_{\sigma,right}$) denote the halfplane bounded by $\ell_{\sigma,left}$ (resp. $\ell_{\sigma,right}$) and containing σ. We use the following simple but important observation (see Figure 3).

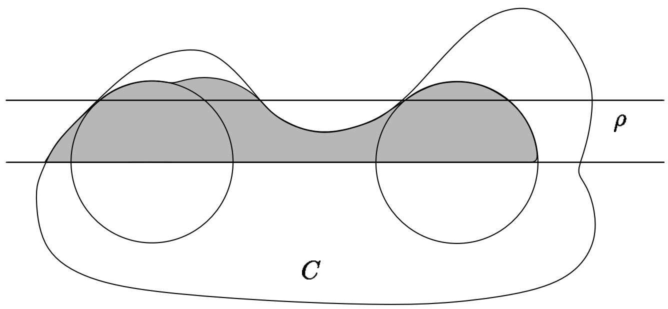

Figure 4: The part C_{bottom} (in gray) and two of its defining disks

Claim 2.1 *Let B be a disk of radius at least 3δ that intersects ρ, and let q be a point on $\rho \cap \partial B$. Let σ be the cell of Σ containing q. Then either $B \cap h^+_{bottom}$ is function defined with respect to ℓ_{bottom}, or $B \cap h^-_{top}$ is function defined with respect to ℓ_{top}, or $B \cap h^+_{\sigma,left}$ is function defined with respect to $\ell_{\sigma,left}$, or $B \cap h^-_{\sigma,right}$ is function defined with respect to $\ell_{\sigma,right}$.*

We next define the set $C_{bottom} \subseteq C$.

$$
\begin{aligned}
C_{bottom} = \bigcup \{ & B' \cap h^+_{bottom} | \\
& B \text{ is a disk of } \mathcal{B}(p) \text{ for some } p \in \rho \cap \partial C, \text{ and} \\
& B' \cap h^+_{bottom} \text{ is function-defined for } \ell_{bottom} \} \ .
\end{aligned}
$$

The set C_{top} is defined analogously. It is easy to see that the sets C_{bottom} and C_{top} are function-defined with respect to ℓ_{bottom} and ℓ_{top}, respectively, and that they have constant description complexity. Let γ be the union of the upper envelope of C_{bottom} and the lower envelope of C_{top}. For each $\sigma \in \Sigma$ we define

$$
\begin{aligned}
C_{\sigma,left} = \bigcup \{ & B' \cap h^+_{\sigma,left} | \\
& B \text{ is a disk of } \mathcal{B}(p) \text{ for some } p \in \sigma \cap \partial C \setminus \gamma, \\
& \text{and } B' \cap h^+_{\sigma,left} \text{ is function-defined for } \ell_{\sigma,left} \} \ .
\end{aligned}
$$

The set $C_{\sigma,right}$ is defined analogously. It is easy to see that the set $C_{\sigma,left}$ (resp. $C_{\sigma,right}$) is function-defined with respect to $\ell_{\sigma,left}$ (resp. $\ell_{\sigma,right}$), that it has constant description complexity, and that its x-span is contained in the x-span of the union of σ and the six cells immediately to its right (resp. left).

Define the function f_{bottom} on ℓ_{bottom} as follows: For $x \in \ell_{bottom}$, if there is no point in C_{bottom} that lies vertically above x, then $f_{bottom}(x)$ is not defined; otherwise, $f_{bottom}(x) = y$, where y is the highest among the points in C_{bottom} that lie vertically above x. The graph of the function f_{bottom} is actually the upper envelope of C_{bottom}. The functions f_{top}, $f_{\sigma,left}$, $f_{\sigma,right}$ are defined analogously. ¿From Claim 2.1 it follows that $\rho \cap \partial C$ can be expressed as the union of the graph of f_{bottom}, the graph of f_{top}, and the union of the graphs of $f_{\sigma,left}$ and $f_{\sigma,right}$ taken over all $\sigma \in \Sigma$.

Next we claim that $C_{\sigma,left}$ and $C_{\sigma,right}$ are not empty only for a constant number of cells σ. Indeed, let p be a point on $\sigma \cap \partial C$ such that there exists a disk $B \in \mathcal{B}(p)$ that is function defined say for the left edge of σ, but not for ℓ_{top} nor ℓ_{bottom}. Analyzing the relative positions of p and of the center of B, we deduce that B fully contains the left edge of σ, or the left edge of the cell σ^l immediately to the left of σ. Therefore, either ∂C has a (locally) leftmost point in σ or in σ^l (i.e., ∂C turns rightwards), or ∂C intersects either ℓ_{bottom} or ℓ_{top} within σ or σ^l. Thus in both cases some event occurs either in σ or in σ^l. However, our assumptions concerning the boundary of C imply that both these types of events may occur only a constant number of times, and therefore the number of non-empty sets of the form $C_{\sigma,left}$ or $C_{\sigma,right}$ is bounded by some constant.

3 Proof of Theorem 1.1

Let $\mathcal{C}$ be a set of n κ-curved objects, and let U denote the union of the objects in $\mathcal{C}$. We prove that the combinatorial complexity of ∂U is $O(\lambda_s(n) \log n)$, for some constant s.

Project the objects in $\mathcal{C}$ on the y-axis, and construct a segment tree $\mathcal{T}$ for these projections. Insert the objects of $\mathcal{C}$ into $\mathcal{T}$ according to their projection on the y-axis. For a node μ of $\mathcal{T}$, let y_μ denote the canonical y-interval that is associated with μ, and let $\mathcal{C}_\mu$ be the canonical subset that is stored at μ. We think of y_μ as the horizontal slab whose top (resp. bottom) defining line passes through the top (resp. bottom) endpoint of the y-interval denoted by y_μ. We also store at μ a second subset $\mathcal{D}_\mu$ which is the union of all canonical subsets stored at (the proper) descendants of μ. It is well known (see e.g. [4]) that

$$\sum_\mu (|\mathcal{C}_\mu| + |\mathcal{D}_\mu|) = O(n \log n) \; .$$

Let U_μ denote the union of the objects in $\mathcal{C}_\mu \cup \mathcal{D}_\mu$ restricted to the slab y_μ. We first prove the following easy (and known) claim.

Claim 3.1 *Let w be a vertex on ∂U that is an intersection point of the boundaries of two objects $C_1, C_2 \in \mathcal{C}$, then there exists a node μ of $\mathcal{T}$ such that w lies in the slab y_μ, and either*

1. *both C_1 and C_2 are in $\mathcal{C}_\mu$, or*
2. *one of them is in $\mathcal{C}_\mu$ and the other is in $\mathcal{D}_\mu$.*

Moreover, w is also a vertex on ∂U_μ.

Proof: The first part (i.e., there exists such a node μ) follows from basic properties of segment trees, since the projection of w on the y-axis lies in both the projection of C_1 and the projection of C_2. The second part is also obvious, since w is a vertex on the boundary of the union of any subset of $\mathcal{C}$ that includes both C_1 and C_2. □

We thus distinguish between two types of vertices on ∂U_μ. A vertex of type I is an intersection point between the boundaries of two objects in $\mathcal{C}_\mu$, and a vertex of type II is an intersection point between the boundaries of an object in $\mathcal{C}_\mu$ and an object in $\mathcal{D}_\mu$. Let u_μ be the number of vertices on ∂U_μ of type I and type II. We prove that $u_\mu = O(\lambda_s(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$, and therefore

$$\sum_\mu u_\mu = \sum_\mu O(\lambda_s(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|)) = O(\lambda_s(n) \log n) \; .$$

This equation together with Claim 3.1 yields the desired result, i.e., the number of vertices on ∂U is $O(\lambda_s(n) \log n)$.

Consider a node μ of $\mathcal{T}$, and let d be the width of the slab y_μ. We partition the slab y_μ into $3/\kappa$ horizontal strips each of width $\frac{\kappa}{3}d$. We partition each of these strips into disjoint squares $\sigma_1, \sigma_2, \ldots$ of edge-length $\frac{\kappa}{3}d$, by adding vertical walls (see Figure 5). Consider any one of the strips ρ, and let l_1 (resp.

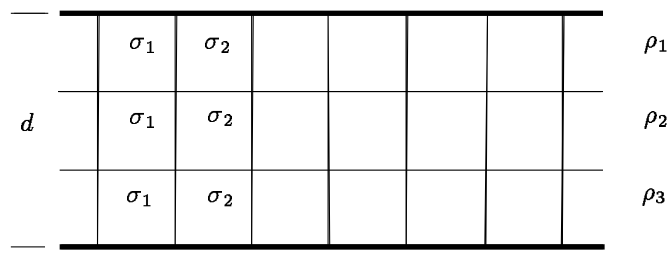

Figure 5: The slab y_μ partitioned into 3 strips ρ_1, ρ_2, ρ_3

l_2) denote its lower (resp. upper) defining line. We show that the number of vertices on ∂U_μ of type I and type II that lie in ρ is $O(\lambda_s(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$, and thus obtain that $u_\mu = O(\lambda_s(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$ (since y_μ was partitioned into a constant number of strips).

Clearly any object in $\mathcal{C}_\mu$ has diameter at least d. Let C be an object in $\mathcal{C}_\mu \cup \mathcal{D}_\mu$ whose diameter is at least d. In Section 2, we proved the following claim.

Claim 3.2 *It is possible to obtain from C a constant number of (not necessarily connected) parts, such that (i) each of the parts is function-defined with respect to either l_1, l_2, or a line containing a vertical wall in ρ, (ii) each of the parts has constant description complexity, (iii) those parts that are function-defined with respect to a line containing a vertical wall e, are contained in a vertical slab defined by a section of ρ that begins at e and is seven squares wide, and (iv) if p is a point on $\rho \cap \partial C$, then p lies on the (appropriate) envelope of one of the parts.*

Let $\mathcal{E}_\mu$ be the set of all objects in $\mathcal{C}_\mu \cup \mathcal{D}_\mu$ with diameter at least d, and let $\mathcal{F}_\mu$ be the set $(\mathcal{C}_\mu \cup \mathcal{D}_\mu) \setminus \mathcal{E}_\mu$. We partition all objects in $\mathcal{E}_\mu$ as in Claim 3.2. Let γ_1 (resp. γ_2) denote the upper envelope (resp. lower envelope) of all parts that are function-defined with respect to l_1 (resp. l_2). The combinatorial complexity of γ_i is $O(\lambda_{s_0}(m_i))$, where $m_i = O(|\mathcal{E}_\mu|) = O(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|)$ is the number of parts that are function-defined with respect to l_i, $i = 1, 2$; [13]. For each square σ, let σ_l (resp. σ_r) denote the right envelope (resp. left envelope) of all parts that are function-defined with respect to the left (resp. right) edge of σ. The combinatorial complexity of σ_l (σ_r) is $O(\lambda_{s_0}(m_l))$ (resp.

$O(\lambda_{s_0}(m_r)))$, where m_l (resp. m_r) is the number of parts that are function-defined with respect to the left (resp. right) edge of σ. From Claim 3.2 we know that

$$\sum_\sigma (m_l + m_r) = O(|\mathcal{E}_\mu|) = O(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|) \ .$$

Consider now the objects in $\mathcal{F}_\mu$, i.e., the objects in $\mathcal{C}_\mu \cup \mathcal{D}_\mu$ with diameter less than d. Each such object intersects only a constant number of squares of ρ. For each square σ, let $\mathcal{F}_\sigma \subseteq \mathcal{F}_\mu$ be the subset of objects that intersect σ; we have $\sum_\sigma |\mathcal{F}_\sigma| = O(|\mathcal{F}_\mu|)$. Recall that our goal is to bound the number of vertices on ∂U_μ of type I and II that lie in ρ. We bound the number of vertices that appear when considering various pairs of envelopes, and various pairs consisting of an envelope and a subset of $\mathcal{F}_\mu$. That is, when considering a pair of envelopes we count the number of intersection points between the envelopes, or, in other words, if $\mathcal{X}$ and $\mathcal{Y}$ are the two underlying sets of parts, then we count the number of *bichromatic* vertices on the boundary of the union of the objects in $\mathcal{X} \cup \mathcal{Y}$, where a vertex is bichromatic if it lies on the boundary of an object of $\mathcal{X}$ and on the boundary of an object of $\mathcal{Y}$. And when considering a pair consisting of an envelope and a subset $\mathcal{F}'$ of $\mathcal{F}_\mu$, we count the number of *bichromatic* vertices on the boundary of the union of the objects in $\mathcal{X} \cup \mathcal{F}'$, where $\mathcal{X}$ is the set of parts underlying the envelope. More precisely, we bound the number of vertices that appear when considering the following pairs.

(a) (γ_1, γ_2)

(b) For each square σ,
(γ_1, σ_l), (γ_1, σ_r), (γ_2, σ_l), (γ_2, σ_r).

(c) For each square σ,
(σ_l, σ_r).

(d) $(\gamma_1, \mathcal{F}_\mu)$, $(\gamma_2, \mathcal{F}_\mu)$

(e) For each square σ,
$(\sigma_l, \mathcal{F}_\sigma)$, $(\sigma_l, \mathcal{F}_\sigma)$.

We now claim that all 'interesting' vertices appear.

Claim 3.3 *If w is a vertex on ∂U_μ of type I or II that lies in ρ, then either (i) w is a vertex of one of the envelopes considered, or (ii) w appears when one of the above pairs is considered.*

Proof: If w is of type I, that is, w is an intersection point of the boundaries of two objects in $\mathcal{C}_\mu$. Then clearly w is either a vertex of one of the envelopes γ_1, γ_2, or σ_l, σ_r, for a square $\sigma \in \Sigma$, or a vertex that appears when considering one of the pairs listed in **(a)**, **(b)**, and **(c)** above.

If w is of type II, that is, w is an intersection point of the boundaries of an object in $\mathcal{C}_\mu$ and an object in $\mathcal{D}_\mu$, then we distinguish between two cases. If the object from $\mathcal{D}_\mu$ is large, i.e., it is in $\mathcal{E}_\mu$, then, as before, w is either a vertex of an envelope or appears when considering a pair of envelopes. Otherwise, the object from $\mathcal{D}_\mu$ is small (i.e., it is in $\mathcal{F}_\mu$), and w is a vertex that appears when considering one of the pairs listed in **(d)** and **(e)**. $\square$

Notice that we also count many 'uninteresting' vertices such as vertices that are formed by two objects in $\mathcal{D}_\mu$, or vertices that 'do not make it' to the boundary of the full union.

In the next section we prove a key lemma stating that if γ is an envelope as those defined above, and $\mathcal{A}$ is a set of κ-curved objects, then the number of 'visible' bichromatic vertices on γ for which the larger object (of the two objects forming the vertex) comes from the set $\mathcal{X}$ underlying γ is $O(\lambda_{s_0}(|\mathcal{X}|)+\lambda_s(|\mathcal{A}|))$, where a visible vertex is a vertex on the boundary of the union of $\mathcal{X} \cup \mathcal{A}$, and the size of a part in $\mathcal{X}$ is the size of the object to which it belongs.

We employ this lemma to bound the number of vertices that appear when considering the pairs listed above. We can immediately apply the lemma to the two pairs of **(d)** above, since each object in the set underlying γ_i is larger than all objects in $\mathcal{F}_\mu$. Thus we obtain an $O(\lambda_s(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$ bound for these two pairs. Similarly, we can apply the lemma to the pairs of **(e)**. Recalling that the total complexity of the envelopes corresponding to vertical walls in ρ is $O(\lambda_{s_0}(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$, and that $\sum_\sigma |F_\sigma| = O(|\mathcal{D}_\mu|)$, we obtain a bound of $O(\lambda_s(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$ for all the pairs of **(e)** together.

In order to apply the lemma to the pairs of **(b)**, we first observe that when a pair (γ_i, σ_z), $i \in \{1,2\}, z \in \{l,r\}$, is considered, we may restrict γ_i to a section of ρ of width seven squares beginning at σ. However, there is still a problem, since it is not true anymore that the larger object (of the two objects forming a countable vertex) always comes from the same underlying set. We thus consider a pair (γ_i, σ_z) twice. First we bound the number of vertices on γ_i for which the smaller object (of the two objects forming it) comes from the set underlying σ_z, by applying the lemma, and then we bound the number of vertices on σ_z for which the smaller object comes from the set underlying γ_i, again by applying the lemma. In this way we bound all vertices that appear when considering a pair (γ_i, σ_z), and obtain a bound of $O(\lambda_s(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$ for all the pairs of **(b)** together.

We can immediately bound the number of vertices that appear when considering the pair of **(a)** or the pairs of **(c)**, and obtain in both cases a bound of $O(\lambda_{s_0}(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$.

We thus conclude that the number of vertices of ∂U_μ of type I or II that lie in ρ is $O(\lambda_s(|\mathcal{C}_\mu| + |\mathcal{D}_\mu|))$, leading as detailed above to the main theorem.

Theorem 3.4 *The combinatorial complexity of ∂U is $O(\lambda_s(n) \log n)$, for some constant s.*

4 Proving the Key Lemma

Let ρ be a strip of width δ and denote by l its bottom boundary. Let $\mathcal{S} = \{S_1, \ldots, S_m\} \subseteq \mathcal{C}$ be a set of m large input objects, that is, the diameter of S_i is at least $3\delta/\kappa$, $i = 1, \ldots, m$. Apply the process described in Section 2 to the objects $S_1, \ldots, S_m$, and let $O_1, \ldots, O_m$ be the m bottom parts that are obtained. Consider γ, the upper envelope of $O_1, \ldots, O_m$, and denote by R the region between γ and l. Let $\mathcal{A} \subseteq \mathcal{C}$ be a set of k input objects. We wish to bound the number of bichromatic vertices on the boundary of $V = R \cup (\cup\mathcal{A})$ that lie on γ and for which the larger of the two appropriate objects comes from $\mathcal{S}$.

We divide each $A \in \mathcal{A}$ into a constant number of primitive objects $\alpha_1, \alpha_2, \ldots$ by vertically decomposing A. That is, for each of the locally x-extreme points p on ∂A, if we remain in A when moving slightly upwards or downwards from p, then we draw a vertical segment beginning at p and directed upwards (alternatively, downwards), until it hits ∂A. Denote by $\mathcal{A}'$ the set of primitive objects that is obtained; $|\mathcal{A}'| = O(k)$. (Notice that a primitive object is trapezoid-like, it is defined by (at most) two vertical walls and by two x-monotone curves, a top curve and a bottom curve.)

When walking along γ from left to right, let L_1 (resp. L_1') be the sequence of names of primitive objects in $\mathcal{A}'$ corresponding to the relevant bichromatic vertices of V that lie on top (resp. bottom) boundaries of primitive objects in $\mathcal{A}'$. In the remainder of this section we prove the following lemma.

Lemma 4.1 (Key Lemma) $|L_1| = |L_1'| = O(\lambda_{s_0}(m) + \lambda_s(k))$*, for some constant* s*.*

Notice that whenever there are more than s_0 consecutive occurrences of the same name, there must be a vertex of γ somewhere in between. Thus, if we replace in the sequence L_1 (resp. L_1') all consecutive occurrences of a name by a single representative occurrence of that name, we remain with a sequence L_2 (resp. L_2'), and $|L_1| = |L_2| + O(\lambda_{s_0}(m))$ (resp. $|L_1'| = |L_2'| + O(\lambda_{s_0}(m))$). In the claim below we prove that $|L_2| = |L_2'| = O(\lambda_s(k))$, for some constant s, and therefore $|L_1| = |L_1'| = O(\lambda_{s_0}(m) + \lambda_s(k))$.

Proposition 4.2 *The sequence* L_2 *(alternatively* L_2'*) is a Davenport-Schinzel sequence [13] of order* $2s_0 + 1/\kappa + c$*, where* c *is a small constant.*

Proof: Consider first the sequence L_2. Assume that there are two primitive objects $\alpha, \beta \in \mathcal{A}'$ with top boundaries $\overline{\alpha}$ and $\overline{\beta}$, respectively, for which there exists a long subsequence of L_2 of the form $\alpha^1\beta^1\alpha^2\beta^2\ldots\alpha^t\beta^t$ (or $\alpha^1\beta^1\alpha^2\beta^2\ldots\alpha^t\beta^t\alpha^{t+1}$). We focus on the x-interval whose endpoints are β^1 and α^t (or more precisely the first occurrence in the sequence of occurrences represented by β^1 and the last occurrence in the sequence represented by α^t). Both top boundaries $\overline{\alpha}$ and $\overline{\beta}$ are defined over the entire interval.

Consider four consecutive representatives β^{2i-1}, α^{2i}, β^{2i}, α^{2i+1}. (If t is even then we disregard the last two representatives β^{t-1} and α^t.) We restrict our attention to the vertical slab ψ whose left bounding line passes through the first occurrence in the sequence of occurrences represented by β^{2i-1}, and whose right bounding line passes through the last occurrence in the sequence represented by α^{2i+1}. Let A and B be the objects of $\mathcal{A}$ from which α and β were obtained.

If α and β intersect within ψ, then we ignore this quadruple, since this implies that the boundaries of A and B intersect within ψ, and therefore there are at most s_0 such quadruples. Thus we assume that either α is above β in ψ, or vice versa. We show that the width of ψ (under this assumption) is at least $2\kappa \cdot \mathit{diam}(C)$, where C is the smaller object among A and B, and therefore there can be at most $\frac{1}{2\kappa}$ such quadruples.

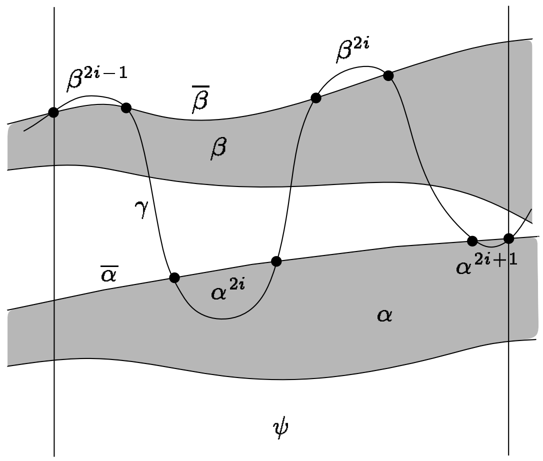

Figure 6: β is above α in ψ

Assume first that β is above α (see Figure 6). We restrict our attention further to the triple $\alpha^{2i}, \beta^{2i}, \alpha^{2i+1}$. Consider p the vertex corresponding to the first occurrence in the sequence represented by β^{2i}, and let $S \in \mathcal{S}$ be the object to which the arc of γ that passes through p belongs. Assume that ∂S exits β at p, and let q be the first point to the right of p on $\overline{\beta}$ where ∂S enters β (see Figure 7, left). (If ∂S enters β at p, then we define q to be the first point to the left of p on $\overline{\beta}$ where ∂S exists β, and proceed similarly.) We now think of α^{2i} as the rightmost intersection point corresponding to it, and of α^{2i+1} as the leftmost intersection point corresponding to it.

We move γ rigidly downwards, varying the points α^{2i}, p, q and α^{2i+1} accordingly, until p and q coincide at a point x on $\overline{\beta}$ (see Figure 7). In other words, during this process, p is the (constantly moving rightwards) exit point of ∂S and q is the (constantly moving leftwards) entrance point of ∂S, α^{2i} is the rightmost intersection point of γ and $\overline{\alpha}$ to the left of p, and α^{2i+1} is the leftmost intersection point of γ and $\overline{\alpha}$ to the right of q. Notice that the path traced by p (alternatively, q) on $\overline{\beta}$ is not necessarily connected (see Figure 7). At the end of this process, ∂S passes through x and lies below $\overline{\beta}$ in a small neighborhood of x. Clearly the final location of α^{2i} is more to the right than the initial location of α^{2i}, and the final location of α^{2i+1} is more to the left than its initial location.

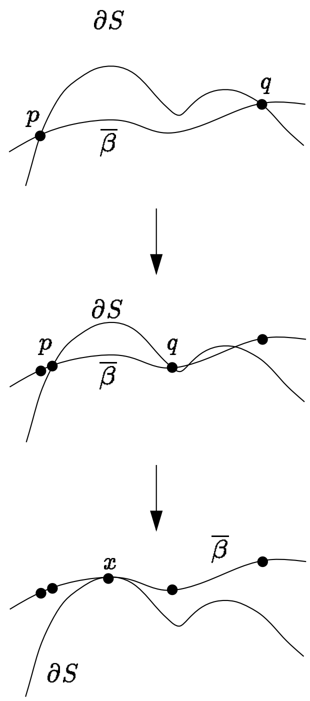

Figure 7: Translating S downwards

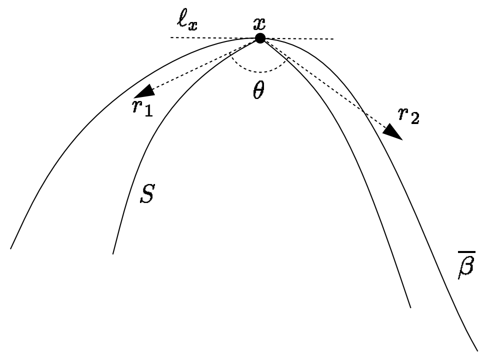

Figure 8: Proof of $D \subseteq S$

If $\overline{\beta}$ does not have a tangent at x, then we may ignore this quadruple, since there are at most s_0 such points on ∂B. Therefore, we assume that $\overline{\beta}$ does have a tangent at x, and let D be the critical disk for the point x on ∂B, i.e., D is a disk of radius $\kappa \cdot diam(B)$, D is contained in B and its boundary passes through x. We now claim that the disk D is also contained in S, and therefore it is contained in the region lying below γ.

Observe that if (as we assume) ∂B has a tangent at x, then so does ∂S. Assume this is false, and let r_1 and r_2 be the two rays tangent from the left and from the right, respectively, to ∂S at x (see Figure 8). S lies locally below both of them. Let θ be the inward angle between r_1 and r_2, and let ℓ_x be the tangent to ∂B at x. If $\theta < \pi$, then it is impossible to draw a disk that is contained in S and whose boundary passes through x. If, on the other hand, $\theta > \pi$, then either r_1 or r_2, say r_2, is above ℓ_x, but then all points of ∂B to the right of x and close enough to x, are below ∂S, which contradicts the way in which S was translated. Thus we conclude that $\theta = \pi$, and ∂S has a tangent at x. Moreover, this tangent is necessarily ℓ_x. Since S is larger than B, D is contained in the (unique) critical disk D' of S for the point x.

The last claim implies that D is contained in $\beta \cap \psi$, since if it is not, then the boundary of D must intersect one of the bounding lines of ψ at two points lying between the bottom and top boundaries of β. But if so γ cannot intersect the top boundary of α within the slab ψ on both sides of x (since $D \subseteq S$).

We now claim that (the current) α^{2i} lies completely to the left of D, and (the current) α^{2i+1} lies completely to the right of D (and therefore this is surely true for the initial α^{2i} and α^{2i+1}). Therefore the horizontal distance between the initial α^{2i} and α^{2i+1} is at least $2\kappa \cdot diam(B)$. The claim is correct since α lies below β in ψ, D is contained in $\beta \cap \psi$ and D is contained in the region below γ, and $\overline{\alpha}, \overline{\beta}$, and γ are x-monotone. If α is above β, then we consider the triple $\beta^{2i-1}, \alpha^{2i}, \beta^{2i+1}$ and treat this case analogously.

Consider now the sequence L'_2. If β is above α we consider the triple $\alpha^{2i}, \beta^{2i}, \alpha^{2i+1}$, and if α is above β we consider the triple $\beta^{2i-1}, \alpha^{2i}, \beta^{2i}$. In both cases, we translate γ until S just touches (locally) the top boundary of the lower object, and essentially continue as for the sequence L_2. We describe in detail the case where β is above α, so we consider the triple $\alpha^{2i}, \beta^{2i}, \alpha^{2i+1}$. Consider p the vertex corresponding to the first occurrence in the sequence represented by β^{2i}, and let $S \in \mathcal{S}$ be the object to which the arc of γ that passes through p belongs. Assume that ∂S enters β at p, and let q be the first point to the right of p on the bottom curve of β where ∂S exits β. (If ∂S exits β at p, then we define q to be the first point to the left of p on the bottom curve of β where ∂S enters β, and proceed similarly.) We now think of α^{2i} as the rightmost intersection point corresponding to it, and of α^{2i+1} as the leftmost intersection point corresponding to it. We translate γ rigidly downwards, until ∂S touches $\overline{\alpha}$ at a point x, to the right of p and to the left of q, and ∂S lies below $\overline{\alpha}$ at a neighborhood of x. As above, if A has a tangent at x, then so does S and the two tangents coincide. We now distinguish between two cases. If $diam(B) < diam(A)$, then at x we may draw a disk of radius $\kappa \cdot diam(B)$ which is surely contained in A and in S. Again we claim that the points α^{2i} and α^{2i+1} are now closer to each other and that they are to the left and to the right of the disk we drew. This means that the horizontal distance between the initial α^{2i} and α^{2i+1} is at least $2\kappa \cdot diam(B)$. If however $diam(B) > diam(A)$, then at x we draw a disk of radius $\kappa \cdot diam(A)$, which is also contained in S since $diam(S) > diam(B)$, and the horizontal distance between the initial α^{2i} and α^{2i+1} is at least $2\kappa \cdot diam(A)$. □

5 Point Location Queries

We can use the tree $\mathcal{T}$ defined in Section 3 as a first step towards a (dynamic) data structure for answering point location queries for the input set $\mathcal{C}$. We concentrate on the problem of finding (efficiently) an object of $\mathcal{C}$ containing a query point q (or reporting that no such object exists). However, data structures for similar types of queries can also be obtained. Please refer to [6] for a description of similar data structures, as well as their applications. Since the techniques are similar, we omit some of the details.

We first present a data structure that supports deletions (only). Fix a node $\mu \in \mathcal{T}$, and let ρ be one of the strips of the slab y_μ. Let l_1 (resp. l_2) be the lower (resp. upper) defining line of ρ. Consider an object $C \in \mathcal{C}$. We define a subset c_1 of ∂C as follows. For each point p on l_1 that lies in C, we draw a vertical segment beginning at p and directed upwards until it hits ∂C. Let c_1 be the union of all upper endpoints of these vertical segments. Clearly the subset $c_1 \subseteq \partial C$ is of constant description complexity. Let $E_{\rho,up}$ be the collection of all subsets of the form c_1 taken over all objects in $\mathcal{C}_\mu$; We define the subset c_2 with respect to the upper line l_2 in an analogus way, and let $E_{\rho,down}$ denote the collection of all subsets of the form c_2 for $C \in \mathcal{C}_\mu$.

Let p be a point on ∂C that is not contained in c_1 nor in c_2, and let σ' be the square in which it lies. We draw a horizontal segment beginning at p and directed into the interior of C, say rightwards, that ends at the *second* right vertical edge of the square c to the right of σ'. The segment fully contained inside C, since C is κ-curved, and $p \in \partial C \setminus c_1 \setminus c_2$. Let $c_{left} \subseteq \partial C$ (resp. $c_{right} \subseteq \partial C$) be the leftmost (resp. rightmost) endoints of these segments. For each vertical edge of a square σ_i we define $E_{\rho,i,left}$ and $E_{\rho,i,right}$ as the collection of all these subsets, taken for each $C \in \mathcal{C}_\mu$ for which c_{left} intesects the square neighboring to the left of σ_i, or c_{right} intesects the square neighboring to the right of σ_i. Observe that the overall complexity of these subsets is $O(n \log n)$.

Claim 5.1 *A point $p \in \rho$ is inside C if the answer to at least on the following conditions, which we call* containment conditions *is true.*

- *p is below the upper envelope of $E_{\rho,\mathrm{up}}$.*
- *p is above the lower envelope of $E_{\rho,\mathrm{down}}$.*
- *p is to the left of the right envelope of $E_{\rho,i,\mathit{left}}$, where σ_i is the square next to the right to the square containing p.*
- *p is to the left of right envelope $E_{\rho,i,\mathrm{right}}$, where σ_i is the square next to the left to the square containing p.*

We construct $\Psi_{\rho,up}$ (resp. $\Psi_{\rho,down}$, $\Psi_{\rho,i,left}$, $\Psi_{\rho,i,right}$), the data structure of [1] for the collections $E_{\rho,up}$ (resp. $E_{\rho,down}$, $E_{\rho,i,left}$, $E_{\rho,i,right}$). This data structre can be constructed in overall time of $O(n^{1+\varepsilon})$, and can check all containment conditions listed in Claim 5.1 in time $O(\log n)$. It can also report all subsets of $E_{\rho,up}$ (resp. $E_{\rho,down}$, $E_{\rho,i,left}$, $E_{\rho,i,right}$) whose graph passes above (resp. below, to the left, to the right) of p in time $O(\log n + k)$, where k is the number of reported subsets. Recall that each such subset corresponds to a different object of $\mathcal{C}_\mu$ containing p. In addition, we can also delete a subset from one of the data structes in time $O(n^\varepsilon)$.

We next construct for each node $\mu \in \mathcal{T}$ and each strip ρ of the slab y_ρ a balance binary search tree $\mathcal{T}_\rho$, whose leaves are associated with $\Psi_{\rho,i,right}$ and $\Psi_{\rho,i,left}$ for all squares σ_i for which $\Psi_{\rho,i,right}$ or $\Psi_{\rho,i,left}$ are not empty. These trees are sorted by the x-coordinate of the squares.

Assume now that a query point q is given, and we want to find some object of $\mathcal{C}$ containing q, or determine that no such object exists. We first query $\mathcal{T}$ to find the set $\mathcal{Z}_q$ of $O(\log n)$ slabs containing q. For each slab $\rho \in \mathcal{Z}_q$ we query $\mathcal{T}_\rho$ to find σ_i, the square of the slabs immediettly to the left of the square containing q. We next query $\Psi_{\rho,up}$, $\Psi_{\rho,down}$, $\Psi_{\rho,i,left}$ and $\Psi_{\rho,i,right}$ to varify all four containmenet conditions of claim 5.1.

Clearly the time needed for a query is $O(\log^2 n)$. Reporting *all* objects containing a quesry points is carried in the same manner, and therefor doable in time $O(\log^2 n + k)$, where k is the number of objects reported.

To perform a deletion of an object C from $\mathcal{C}$, we find all data structures containing C, and delete C from each one. This is carried out as described in [1], in time $O(n^\varepsilon)$ for each node $\mu \in \mathcal{T}$. Since the number of such nodes is $O(\log n)$, the overall running time is $O(\log n \cdot n^\varepsilon) = O(n^\varepsilon)$. Instead of perfroming a deletion from $\mathcal{T}$ (which is impossibe in a standard segment tree), we construct the data structure from scratch each time that the number of deleted objects is more than the number of non-deleted objects. Clearly the amortized running time for a deletion is still $O(n^\varepsilon)$. Using known techniques, this can also be the worst-case time bound.

If in addition to deletion, we want to support insertion of object to $\mathcal{C}$, we use the decomposition technique of [11]. Using this idea, an insertion is doable in time $O(n^\varepsilon \cdot \log n) = O(n^\varepsilon)$, the (asymptotic) time for a deletion operation remains unchanged, and the query time increases by a multiple factor of $\log_2 n$, so it is doable in time $O(\log^3 n)$. To summarize

Theorem 5.2 *Let C be a set of n convex κ-curved objects in the plane. We can preprocess C in time $O(n^{1+\varepsilon})$, into a data structure of size $O(n^{1+\varepsilon})$, such that finding an object of C containing a query point q can be done in time $O(\log^3 n)$. Moreover, we can insert or delete an object into/from C in time $O(n^{1+\varepsilon})$. In addition, we can report all k objects containing a query point in time $O(\log^3 n + k)$.*

6 Conclusion

We have proven that the combinatorial complexity of the boundary of the union of a set of n κ-curved objects is $O(\lambda_s(n) \log n)$, for some constant s. This bound improves the recent bound of Efrat and Sharir [5] for the case of convex κ-curved objects (e.g., fat ellipses). (They obtained a bound of $O(n^{1+\varepsilon})$ for convex fat objects.) This bound is also the first non-trivial bound for the case of non-convex κ-curved objects (e.g., rounded heart-shaped objects).

References

[1] P.K. Agarwal, A. Efrat and M. Sharir, Vertical decomposition of shallow levels in 3-dimensional arrangements and its applications, *Proc. 11th ACM Symp. Comput. Geom.*, 1995, 39–50.

[2] P.K. Agarwal, M.J. Katz and M. Sharir, Computing depth orders and related problems, *Comp. Geom. Theory and Appls* 5 (1995), 187–206.

[3] M. de Berg, Linear size binary space partitions for fat objects, *Proc. 3rd Annu. European Symp. Algorithms*, 1995, 252–263.

[4] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, Algorithms for bichromatic line segment problems and polyhedral terrains, *Algorithmica* 11 (1994), 116–132.

[5] A. Efrat and M. Sharir, On the complexity of the union of fat objects in the plane, *Proc. 13th ACM Symp. on Computational Geometry*, 1997, 104–112.

[6] A. Efrat, M.J. Katz, F. Nielsen and M. Sharir, Dynamic Data Structures for Fat Objects and their Applications, *Proc. 5th Workshop Algorithms and Data Struct.*, 1997, 297–306.

[7] M.J. Katz, 3-D vertical ray shooting and 2-D point enclosure, range searching, and arc shooting, amidst convex fat objects, *Comp. Geom. Theory and Appls* 8 (1997), 299–316.

[8] M.J. Katz, M.H. Overmars and M. Sharir, Efficient hidden surface removal for objects with small union size, *Comp. Geom. Theory and Appls* 2 (1992), 223–234.

[9] M. van Kreveld, On fat partitioning, fat covering, and the union size of polygons, *Proc. 3rd Workshop Algorithms Data Struct.*, 1993, 452–463.

[10] J. Matoušek, J. Pach, M. Sharir, S. Sifrony and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.* 23 (1994), 154–169.

[11] M.H. Overmars and J. van Leeuwen, Two general methods for dynamizing decomposable searching problems, *Computing* 26 (1981), 155–166.

[12] M.H. Overmars and A.F. van der Stappen, Range searching and point location among fat objects, *J. Algorithms* 21 (1996), 629–656.

[13] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[14] F. van der Stappen, *Motion planning amidst fat obstacles*, Ph.D. Dissertation, Utrecht University, 1994.

[15] J. Vleugels, *On fatness and fitness — realistic input models for geometric algorithms*, Ph.D. Dissertation, Utrecht University, 1997.