

Computational and Structural Advantages of Circular Boundary Representation^{*}

O. Aichholzer¹, F. Aurenhammer¹, T. Hackl¹,
B. Jüttler², M. Oberneder², and Z. Šir²

¹ University of Technology Graz, Austria
{oach,thackl}@ist.tugraz.at, auren@igi.tugraz.at

² Johannes Kepler University of Linz, Austria
{bert.juettler,margot.oberneder,zbynek.sir}@jku.at

Abstract. Boundary approximation of planar shapes by circular arcs has quantitative and qualitative advantages compared to using straight-line segments. We demonstrate this by way of three basic and frequent computations on shapes – convex hull, decomposition, and medial axis. In particular, we propose a novel medial axis algorithm that beats existing methods in simplicity and practicality, and at the same time guarantees convergence to the medial axis of the original shape.

1 Introduction

The plain majority of algorithms in computational geometry have been designed for processing *linear* objects, like lines, planes, or polygons. On the one hand, this is certainly due to the fact that many interesting and deep computational and combinatorial questions do arise already for inputs of this simple form. Again, the pragmatic reason is that algorithms for linear objects are usually both easier to develop and simpler to implement. To make things work for nonlinear objects, which arise frequently in practical settings, such objects are usually approximated in a piecewise-linear manner and up to a tolerable error. Existing approaches [10] to directly extending polygonal algorithms to curved objects are rare and, due to their generality, of limited practical use.

In its simplest form, the input object is a single planar shape, \mathcal{A} , with curved and connected boundary $\partial\mathcal{A}$. Frequent tasks to be performed on \mathcal{A} – each being prior to a variety of more involved computations – include constructing the convex hull of \mathcal{A} , decomposing \mathcal{A} into primitives, and calculating the medial axis of \mathcal{A} . These tasks are well investigated in the case of polygonal shapes. In certain situations, however, the number of line segments required for approximating $\partial\mathcal{A}$ with high accuracy may be prohibitively large. Even more seriously, making a piecewise-linear approximation of $\partial\mathcal{A}$ and invoking a polygonal-shape algorithm may generate results that are topologically incorrect; the medial axis is a well-known example.

^{*} Supported by the Austrian FWF JRP 'Industrial Geometry', S9200.

The intention of the present paper is to highlight the use of circular arcs for boundary representation. It is well known that for nonlinear curve segments the approximation order increases in comparison to using straight-line segments. In particular, if a given accuracy ε is achieved by using N line segments, then as few as $n = \Theta(N^{2/3})$ circular arcs can accomplish the same. This has been an issue in approximation theory, but in computational geometry this gain seems to have been less valued than eliminating small factors in the complexity of the subsequently applied algorithm. Boundary approximation by circular arcs may be of advantage also in a qualitative respect. For instance, it avoids the mentioned topological inconsistencies in medial axis computations, and it supports the computation of shape offsets, as the class of shapes bounded by circular arcs is closed under offset operations.

We will show that for the three basic problems mentioned above – convex hull, triangulation, and medial axis – simple and practical, though still efficient algorithms exist that work for circular arc inputs. The first two problems are less demanding; we treat them mainly to point out the respective favorable (in our opinion) approach, whose practicality shall encourage the use of circular arc boundary representation. Nevertheless, substantial differences to the polygonal case occur; see below. For computing the medial axis, we propose a novel and extremely simple algorithm that is based on a known (though less recognized) decomposition lemma. After having computed a purely combinatorial description of the medial axis using tailored shape splitting, its individual parts (conics and line segments, like in the polygonal case) are reassembled *without* the need of merging.

Suitable circular arc approximations of shapes can be found in linear time. In summary, the obtained shape processing algorithms are superior in runtime and output volume to their line segment based counterparts, retain much (if not all) of their simplicity, and are even more natural in some cases.

1.1 Outline and Background

We briefly describe the contributions of this paper and relate them to existing literature.

Section 2 deals with approximating general curves by suitable primitives. This is a topic of importance in geometric modeling and in CAD and NC applications, and many quite recent results are available [24,25,27,33,17,31]. Our aim is to approximate a parametric curve $\mathbf{c}(t)$ by circular arcs. We assume that $\mathbf{c}(t)$ is piecewise-polynomial of constant degree, and we use biarcs (pairs of smoothly joined circular arcs) [30,25,31] as primitives. A straight-forward bisection algorithm for biarc generation already fits our purposes. It uniquely assigns biarcs to parameter intervals, which facilitates the error evaluation. An approximating spline curve \mathbf{b} of size n is computed in $O(n)$ time. It fits the input curve $\mathbf{c}(t)$ in slope at biarc endpoints, and can be tuned to match $\mathbf{c}(t)$ in curvature at certain points (a fact being important in subsequent medial axis computations). Though not being optimal in the number of arcs, the approximation order of \mathbf{b} is still

three [24,31]. In contrast, with line segments one cannot exceed order two, and a polyline of size $N = \Theta(n^{3/2})$ is needed to arrive at the same precision.

The remaining sections propose algorithms for *circular arc shapes* \mathcal{A} , where the boundary $\partial\mathcal{A}$ of \mathcal{A} is given as a simple curve composed of n circular arcs. Choice is guided by efficiency as well as by reducibility to basic operations that have robust implementations [11]. Due to lack of space, we had to skip two sections from this version of the paper. Let us nevertheless provide here a short description of the material they contain.

The first topic is computing the convex hull of a circular arc shape \mathcal{A} . This task is one of the most basic to be performed for a given shape, and has a variety of applications including shape fitting, motion planning, shape separation, and many others. At least four linear-time algorithms have been developed for polygonal shapes [4,16,23,26]. The incremental method by Melkman [26] stands out by its simplicity, and it is this candidate we generalize for circular arc shapes. Compared to the original setting, two difficulties arise. Deciding inclusion for a currently inserted arc in the convex hull constructed so far is no trivial test, and the convex hull cannot be described by a sequence of input vertices of the shape. We show that a runtime of $O(n)$ is still possible. The basic subroutine of the algorithm computes the convex hull of only two circular arcs.

The second topic is shape triangulation, a fundamental building block in algorithms for decomposition, shortest path finding, and visibility – to name a few. Most existing algorithms are meant for polygonal shapes. They partition a given (simple) n -vertex polygon into triangles without introducing Steiner points. Efficient candidates are [14,22,3,18,7] which all show an $O(n \log n)$ runtime. Theoretically more efficient methods do exist, but when aiming at simplicity, choice should be made from the list above. When trying to generalize to shapes \mathcal{A} bounded by circular arcs, we face two problems. First of all, if the use of Steiner points is disallowed, then a partition of \mathcal{A} into primitives bounded by a constant number of circular arcs need not exist. Also, not all triangulation methods are suited to generalization. This applies, for instance, to the extremely simple ear cutting method in [20] which runs in time $O(r \cdot n)$, where r is the number of reflex vertices of \mathcal{A} . The triangulation algorithm we propose is closest to Chazelle's [7]. It manages with an (almost) worst-case minimal number of Steiner points on $\partial\mathcal{A}$, runs in $O(n \log n)$ time, and uses a dictionary as its only nontrivial data structure. The produced primitives are arc triangles with at least one straight edge. The most complex geometric operation is intersecting a circle with a line.

Section 3 is devoted to the medial axis, a frequently used structure associated with a given input shape. Its main applications include shape recognition, solid modeling, pocket machining, and others. Interest in mathematical properties of the medial axis for general shapes found renewal in recent years [9,28,5,6,2]. In our case, where the shape \mathcal{A} is simply connected and $\partial\mathcal{A}$ consists of n circular arcs, its medial axis $M(\mathcal{A})$ is known to be a tree composed of $O(n)$ conic edges. Algorithmic work on the medial axis either concentrated on the case where \mathcal{A} is a polygon [21,7,8], or on general sets of curved arcs [33,19,28,1] (and their Voronoi

diagram) without, however, exploiting the fact that the input arcs define a simple curve. (The various existing methods for computing digital versions of the medial axis are not considered here.) Though theoretically efficient as $O(n \log n)$ or better, these algorithms suffer from involved merge or insertion steps which, even for straight arcs as input, are difficult to implement. In addition, numerical stability issues arise heavily; intersections of conics have to be determined repeatedly which, when not calculated exactly, are bound to cumulate the error.

We present a simple randomized divide-and-conquer algorithm for computing $M(\mathcal{A})$ that overcomes these drawbacks. In contrast to comparable algorithms, the costly part is delegated to the divide step. The basic subroutine there is an inclusion test for an arc in a circle. The merge step is trivial: it concatenates two medial axes. The expected runtime is bounded by $O(n^{3/2})$, but is provably better for most types of shape. For example, $O(n \log n)$ expected time suffices if the diameter of $M(\mathcal{A})$ is $\Theta(n)$. No nontrivial data structures are used.

To guarantee applicability of our methods to approximating the medial axes of general shapes \mathcal{B} , a convergence result is needed. We prove in Section 4 that, for a suitable approximation of $\partial\mathcal{B}$ by biarcs, $M(\mathcal{B})$ is the limit of $M(\mathcal{A})$ when the approximating arc shape \mathcal{A} converges to \mathcal{B} . Related results exist, but either presuppose C^2 conditions on $\partial\mathcal{A}$ not attainable by circular arcs [6], or concern subsets of the medial axis [5] that survive after pruning the Voronoi diagram of point samples from $\partial\mathcal{B}$. (As a negative side effect, the medial axis approximation obtained from a point sample is not C^1 .) It is well known [2] that medial axis convergence is *not* given for polygonal approximations of \mathcal{B} . In conclusion, circular arcs are the simplest possible tool for boundary conversion that guarantees a stable medial axis approximation.

2 Circular Arcs

In order to represent a general shape \mathcal{A} in a form suitable for geometric computations, we discuss methods for approximating $\partial\mathcal{A}$ by circular arcs. We assume that $\partial\mathcal{A}$ is given as a polynomial spline curve of constant degree. Attention is restricted to degree 3, as every free-form curve can efficiently and with any desired precision be converted into cubics [29], and in many applications the input will already be available in this common form [12].

Several approaches to generating circular arc splines exist; see [24] for a review. We consider a simple bisection algorithm consisting of two steps, approximation and error measurement. A geometric primitive b (an arc or a biarc) is fitted to a segment s of the given cubic curve $\mathbf{c}(t)$, and the distance from b to s is computed. The algorithm is relatively easy to implement and still adapts the degrees of freedom to the input data. As a slight disadvantage, the number of primitives (the resulting data volume) is minimal only in the asymptotic sense.

Define the one-sided Hausdorff distance from a primitive b to a segment $s \subseteq \mathbf{c}(t)$ as $\delta(b, s) = \max_{p \in b} \min_{q \in s} \|p - q\|$. (We consider b and s as closed sets.) Let ε denote the error tolerance to be met by the algorithm.

Algorithm BISECT(t_0, t_1)

```

Construct  $b$ 
Compute  $\delta = \delta(b, \mathbf{c}[t_0, t_1])$ 
If  $\delta \leq \varepsilon$  then return  $\{b\}$ 
Else return BISECT( $t_0, \frac{t_0+t_1}{2}$ )  $\cup$  BISECT( $\frac{t_0+t_1}{2}, t_1$ )
    
```

Depending on the primitive b used, Algorithm BISECT produces splines of different quality: merely continuous (C^0) circular arc splines, or tangent continuous (C^1) arc splines. When being content with the former type, we can simply choose for b the unique circular arc passing through the three points $\mathbf{c}(t_0)$, $\mathbf{c}(\frac{t_0+t_1}{2})$, and $\mathbf{c}(t_1)$. To obtain C^1 arc splines, so-called biarcs [30] are utilized.

A *biarc* b consists of two circular arcs with common unit tangent vector at their joint. Usually, b is described by its source x with associated unit tangent vector v_x , and its target y with unit tangent vector v_y . Given these data, there exists a one-parameter family of interpolating biarcs. All possible joints are located on the circle σ passing through x and y and having the same oriented angles with v_x and v_y . Several ways for choosing the joint m have been proposed; see e.g. [25,31]. For many applications, taking $m = \sigma \cap \mathbf{c}[t_0, t_1]$ is appropriate. To calculate m , a polynomial of degree 4 has to be solved (where a closed-form solution is still available). The output is a C^1 arc spline with all arc endpoints sitting on $\mathbf{c}(t)$.

In view of subsequent stable medial axis computations, the choice of m has to be made more carefully. Define an *apex* of $\mathbf{c}(t)$ as a local curvature maximum. The apices split the curve $\mathbf{c}(t)$ into pieces of monotonic signed curvature, so-called *spirals*. Following [25], we aim at approximating spirals of $\mathbf{c}(t)$ by circular arc spirals. To this end, we split $\mathbf{c}(t)$ at its apices. These points can be found by solving polynomials of degree 5. Now, we exploit that spiral biarcs can be constructed that connect two given points x and y , match unit tangents there, and assume a predefined curvature in one of them. Let k_x and k_y be the curvature of $\mathbf{c}(t)$ at x and y , respectively, and suppose $k_x < k_y$. To match curvature at x , we choose the radius of the first arc, b_1 , equal to $r_x = 1/k_x$. The joint m is obtained by intersecting the circle supporting b_1 with the joint circle σ . According to [25], the radii and curvatures satisfy $r_x > r_y > 1/k_y$. When starting the next biarc from y with $r_y = 1/k_y$ (unless y is an apex), monotonicity of signed curvature will be preserved.

Each arc is found in $O(1)$ time, where the constant depends on the degree of the polynomial to be solved. Concerning the error measurement, each produced circular arc b_i has to be matched to its corresponding segment s on the curve \mathbf{c} . We then compute an upper bound for the one-sided Hausdorff distance $\delta(b_i, s)$ by substituting the parametric representation of s into the implicit equation of the circle supporting b_i . A simpler upper bound can be calculated (without polynomial solving) by using Bernstein-Bézier representations [13]. In summary, when algorithm BISECT spans a binary recursion tree with n leaves (the returned n primitives), any of the described arc splines can be constructed in $O(n)$ time.

Let us discuss the asymptotic behaviour of the number n for decreasing tolerance ε . For a given curve $\mathbf{c}(t)$ with domain $[t_0, t_1]$, which is assumed to contain

neither inflections nor apices, we consider primitives having approximation order k . Adapting the analysis in [24,31], we get $\delta = \Theta(h^k)$ for the one-sided Hausdorff distance δ , provided that $\mathbf{c}(t)$ is approximated with (small) parameter step size h , and that k is considered a constant.

This relation implies a general lower bound. For *any* approximation of $\mathbf{c}(t)$ by n primitives of order k , the largest step size satisfies $\Delta t \geq \frac{t_1 - t_0}{n}$. Thus from $\varepsilon \geq \delta$, which is to be achieved by the approximation, and from $\delta = \Theta((\Delta t)^k)$, we get $n = \Omega(1/\varepsilon^{1/k})$. On the other hand, the smallest step size Δt taken by algorithm BISECT satisfies $\Delta t \leq \frac{t_1 - t_0}{n}$. When doubling the step size we have $\delta = \Theta((2\Delta t)^k)$ but $\varepsilon < \delta$, as the tolerance is not yet achieved. Thus $n = O(1/\varepsilon^{1/k})$.

In conclusion, for sufficiently small tolerance ε , the number n of primitives constructed by algorithm BISECT is asymptotically optimal. This is also true in the general case where $\mathbf{c}(t)$ contains inflections and apices, because the resulting number of spirals of $\mathbf{c}(t)$ is independent of n . In conclusion, to arrive at tolerance ε , Algorithm BISECT needs $n = \Theta(1/\sqrt[3]{\varepsilon})$ circular arcs (order 3), whereas $N = \Theta(1/\sqrt{\varepsilon})$ line segments (order 2) have to be invested by any polygonal approximation method.

Lemma 1. *Compared to approximating $\mathbf{c}(t)$ with a polyline, the data volume drops from N to $\Theta(N^{2/3})$ when circular arc splines are used.*

It should be observed that, the other way round, when approximating $\mathbf{c}(t)$ with a point sample (as commonly done for medial axis computations [2]), the data volume increases to $\Theta(n^3)$ compared to n circular arcs.

3 Medial Axis

Let \mathcal{A} be the circular arc shape under consideration. (All objects are considered to be closed sets in the sequel). Call a disk $D \subseteq \mathcal{A}$ *maximal* if there exists no disk D' different from D such that $D \subset D'$ and $D' \subseteq \mathcal{A}$. The medial axis, $M(\mathcal{A})$, of \mathcal{A} is defined as the set of all centers of maximal disks.

As the boundary of \mathcal{A} is a connected and simple curve with n circular arcs, $M(\mathcal{A})$ is finite, connected, and cycle-free [9] and thus forms a tree. $M(\mathcal{A})$ can be decomposed into $O(n)$ *edges*, which are maximal pieces of straight lines and (possibly all four types of) conics. Endpoints of edges will be called *vertices* of $M(\mathcal{A})$.

The contribution of this section is a simple and practical randomized algorithm for computing $M(\mathcal{A})$. It works by divide-and-conquer and accepts as input any description of $\partial\mathcal{A}$ by circular arcs and/or line segments. The costly part is delegated to the divide step, which basically consists of inclusion tests for arcs in circles. The merge step is trivial; it just concatenates two partial medial axes. The expected runtime is bounded by $O(n^{3/2})$, and will be proved to be $O(n \text{ polylog } n)$ for several types of shape. A qualitative difference to existing medial axis algorithms is that a *combinatorial* description of $M(\mathcal{A})$ is extracted first, which can then be directly (and robustly) converted into a geometric representation. We base our algorithm on the following simple though elegant decomposition lemma [9].




Fig. 1. Walk (dashed) and cut (dotted)

Lemma 2. Consider any maximal disk D for \mathcal{A} . Let A_1, \dots, A_t be the connected components of $\mathcal{A} \setminus D$, and denote with p the center of D . The following holds.

$$(1) \quad M(\mathcal{A}) = \bigcup_{i=1}^t M(A_i \cup D) \qquad (2) \quad \{p\} = \bigcap_{i=1}^t M(A_i \cup D)$$

In plain words, having at hands some maximal disk one can compute the medial axes for the resulting components recursively, and then glue them together at a single point. However, the desired efficiency of this strategy calls for a balanced decomposition. Its existence is given below.

Lemma 3. There exists a maximal disk D for \mathcal{A} such that at most $\frac{n}{2}$ arcs from $\partial\mathcal{A}$ are (completely) contained in each component of $\mathcal{A} \setminus D$.

Proof. Each point $p \in M(\mathcal{A})$ corresponds to a unique maximal disk D_p for \mathcal{A} . Let $f(D_p)$ be the number of arcs from $\partial\mathcal{A}$ in the largest component induced by D_p . As long as $f(D_p) > \frac{n}{2}$, the component that realizes $f(D_p)$ is unique, and we can decrease $f(D_p)$ by continuously moving p on $M(\mathcal{A})$ such that D_p enters into this component. This process terminates at some point p^* where $f(D_{p^*}) \leq \frac{n}{2}$. We never move back the way we came, as the component we move out never exceeds a size of $\frac{n}{2}$.

We are left with the algorithmic problem of finding some maximal disk that yields a well-balanced partition. Observe that the optimal point p^* above may be not unique, because the number $f(D_p)$ is invariant under motion of p within the relative interior of any fixed edge $e \subset M(\mathcal{A})$. Let us define $Walk(e)$ as the path length in $M(\mathcal{A})$ from e to p^* . Further, define $Cut(e)$ as the size of the smaller one among the two subtrees which constitute $M(\mathcal{A}) \setminus \{e\}$. See Figure 1. Any tree with small 'cuts' tends to have short 'walks', in the following respect.

Lemma 4. Let e be an edge of $M(\mathcal{A})$, chosen uniformly at random. Then $E[Walk(e)] = \Theta(E[Cut(e)])$.

Proof. Orient all the paths in $M(\mathcal{A})$ away from the point p^* . This defines a partial order \prec on the edges of $M(\mathcal{A})$. We have the set equality

$$\bigcup_{e \in M(\mathcal{A})} \{(a, e) \mid a \prec e\} = \bigcup_{e \in M(\mathcal{A})} \{(e, b) \mid b \succ e\}$$

because either set contains each pair of the relation exactly once. The (disjoint) subsets united in the left set, L , represent all the paths in $M(\mathcal{A})$ between its edges e and p^* . Thus we have $E[Walk(e)] = \frac{1}{m} \cdot |L|$, where m is the number of edges of $M(\mathcal{A})$. The (disjoint) subsets united in the right set, R , represent those subtrees defined by the edges e of $M(\mathcal{A})$ which avoid p^* . If we neglect subtrees of sizes larger than $\frac{m}{2}$, then the cardinality of the set drops by a constant factor (of at most 4, if \prec would be a total order, hence less). This implies $\frac{1}{m} \cdot |R| > E[Cut(e)] > \frac{1}{m} \cdot \frac{|R|}{4}$. The lemma now follows from $|R| = |L|$.

Lemma 4 motivates the following disk finding algorithm which combines random cutting with local walking. Its main subroutine, $MAX(b)$, selects for an arc $b \subset \partial\mathcal{A}$ its midpoint x and returns the unique maximal disk for \mathcal{A} with x on its boundary. For the ease of description, we assume that this disk splits \mathcal{A} into exactly two components. Let $c \geq 3$ be a (small) integer constant.

Procedure CUT(\mathcal{A})

Put $A' = \mathcal{A}$
 Repeat
 Choose a random arc b of $\partial A'$
 Compute $D=MAX(b)$ and let \mathcal{A}_0 be the larger component of \mathcal{A} induced by D
 Assign $A' = A' \cap \mathcal{A}_0$
 Until \mathcal{A}_0 contains less than $n - \frac{n}{c}$ arcs
 Report D

Procedure WALK(\mathcal{A})

Choose a random arc b of $\partial\mathcal{A}$
 Compute $D=MAX(b)$
 Put \mathcal{A}_0 =larger component induced by D
 While \mathcal{A}_0 contains $> n - \frac{n}{c}$ arcs do
 Let b_1 (b_2) be the first (last) complete arc of $\partial\mathcal{A}$ in \mathcal{A}_0
 Find $D_1=MAX(b_1)$ and $D_2=MAX(b_2)$
 Put $\mathcal{A}_0 =$ smallest of the respective larger components of \mathcal{A} for D_1 and D_2
 Let $D \in \{D_1, D_2\}$ be the respective disk
 Report D

The disk finding algorithm now runs $CUT(\mathcal{A})$ and $WALK(\mathcal{A})$ in parallel and terminates as soon as the first disk is reported. To analyze its runtime, let us first consider the assignment of arcs on $\partial\mathcal{A}$ to edges of $M(\mathcal{A})$, as done in subroutine MAX . Namely, if $MAX(b)=D$ then arc b is mapped to the edge e that contains the center of D . Observe that either 0, 1, or 2 arcs are mapped to a fixed edge. Moreover, no two unaddressed edges and no two doubly addressed edges are neighbored. This assignment is sufficiently uniform to convey randomness from arcs to edges. So Lemma 4 applies, and in the worst case of walk length being balanced with cut number, a bound of $O(\sqrt{n})$ on the expected number of loop executions in at least one of $CUT(\mathcal{A})$ and $WALK(\mathcal{A})$ holds.

The costly part in both procedures is their subroutine MAX , whose expected number of calls obeys the same bound. $D=MAX(b)$ has a simple implementation which runs in $O(n)$ time: We initialize the disk D as the (appropriately oriented)

halfplane that supports b at its midpoint x and, for all remaining arcs $b_i \subset \partial\mathcal{A}$ that intersect D , we shrink D so as to touch b_i while still being tangent to b at x . The most complex operation needed for shrinking D is an inclusion test of a point in a circle. In particular, and unlike previous medial axis algorithms, no conics take part in geometric operations.

In summary, the randomized complexity for computing the medial axis is given by $T(n) = T(\frac{1}{c}n) + T((1 - \frac{1}{c})n) + O(n^{3/2}) = O(n^{3/2})$. In many cases, however, will the algorithm perform substantially better. Let d be the graph diameter of $M(\mathcal{A})$. Then the loop in $\text{WALK}(\mathcal{A})$ is executed less than d times. So, for example, if $d = \Theta(\log n)$ then an overall runtime of $O(n \log^2 n)$ is met. For the other extreme case, $d = \Theta(n)$, our strategy is even faster. With constant probability, an edge on the diameter is chosen, and $\Theta(n)$ such edges e have $\text{Cut}(e) = \Theta(n)$. The expected number of loop executions in $\text{CUT}(\mathcal{A})$ now is only $O(1)$, and an $O(n \log n)$ algorithm results. We conjecture that the latter situation is quite relevant in practice. In many applications, for typical shapes their medial axes will not branch extensively. Even if so, the branching will be independent of n , because each branch will be approximated by a large number of circular arcs in order to achieve the necessary precision.

The output of the algorithm is a list of $O(n)$ points on $M(\mathcal{A})$, namely, the centers of the splitting disks, plus a list of $O(n)$ edges connecting them. Each edge is given implicitly by its defining two arcs on $\partial\mathcal{A}$. To make sure that the reported point list includes all the vertices of $M(\mathcal{A})$, base cases that involve constantly many (pieces of) original arcs from $\partial\mathcal{A}$ have to be solved directly. (The constant equals 2 or 3 if $\partial\mathcal{A}$ is C^1 .) Note that the algorithm works exclusively on $\partial\mathcal{A}$ except for a final step, where the conic edges of $M(\mathcal{A})$ are explicitly calculated and reassembled. This gives rise to increased numeric stability in comparison to existing approaches.

Opposed to approximating $\partial\mathcal{A}$ with the same accuracy by a polyline of size N , our circular arc algorithm takes $O(n^{3/2}) = O(N)$ time; see Lemma 1 in Section 2. Thus, even for (probably rare) worst-case inputs, our simple algorithm competes asymptotically well with previous methods. Another advantage over polygonal approximation is described in Section 4.

4 Convergence

A well-known unpleasant phenomenon of the medial axis is its instability under perturbations of the shape boundary. Several papers discussing this issue have been published recently. A result in [6] shows that stability is, in general, not given unless perturbations are C^2 . To deal with general shapes, the so-called λ -medial axis has been introduced as a tool in [5]. After drawing a point sample from the shape boundary, the Voronoi diagram of these points is constructed and pruned appropriately. The λ -medial axis converges to the original for vanishing sample distance. Drawbacks are the large sample size for a close (and homotopy-equivalent) approximation, the lack of its C^1 behavior, and the need of computing a general planar Voronoi diagram.

We prove in this section that medial axis convergence under the Hausdorff distance comes as a byproduct of the careful (though, of course, still C^1) biarc boundary conversion described in Section 2. We start with two technical lemmas, whose proofs are omitted due to lack of space.

For some shape \mathcal{A} and a point $p \in M(\mathcal{A})$, let D_p denote the unique maximal disk with center p . Recall that $M(\mathcal{A})$ is defined as the union of the centers of all maximal disks. Define $\xi_p \leq \pi$ as the largest angle at p spanned by two points in the set $D_p \cap \partial\mathcal{A}$. The assertion below does not assume any regularity condition for the shape boundaries.

Lemma 5. *Let \mathcal{A} and \mathcal{B} be two shapes whose (two-sided) Hausdorff distance satisfies $H(\partial\mathcal{A}, \partial\mathcal{B}) = \varepsilon$. Define $k = \frac{4}{1 - \cos(\xi_p/2)}$ and let D_p denote any maximal disk for \mathcal{A} whose radius r fulfills $r > k \cdot \varepsilon > 0$. Then there exists a maximal disk D_q for \mathcal{B} such that $\|p - q\| < k \cdot \varepsilon$.*

Define a *leaf* of the medial axis as a vertex with a single incident edge. The following lemma describes the behavior of $M(\mathcal{A})$ in the vicinity of its leaves. Recall that an apex of $\partial\mathcal{A}$ is a point of maximal curvature.

Lemma 6. *For an apex x of $\partial\mathcal{A}$, consider the unique maximal disk D_p that avoids a segment of $\partial\mathcal{A}$ through x with fixed (small) length ℓ . Further, consider the maximal disk D_q osculating at x . If $\partial\mathcal{A}$ is piecewise analytic C^2 in the neighborhood of x then*

$$\|q - p\| \rightarrow 0 \text{ as } O(\ell^2).$$

We are now prepared to prove the claimed convergence result. Slightly more general than in Section 2, we assume that $\partial\mathcal{A}$ for the original shape \mathcal{A} is C^2 and piecewise analytic. (These requirements are fulfilled if $\partial\mathcal{A}$ is a cubic spline.) The proof generalizes easily to the case where $\partial\mathcal{A}$ is an arbitrary concatenation of analytic pieces, and thus, in particular, is allowed to contain 'sharp' vertices.

Let \mathcal{B}_n denote some circular arc shape that comes from approximating $\partial\mathcal{A}$ by a suitable biarc spline; see Section 2. For sufficiently large n , each leaf of $M(\mathcal{A})$ is also a leaf of $M(\mathcal{B}_n)$, and all leaves of $M(\mathcal{B}_n)$ are contained in $M(\mathcal{A})$. This is because the spline preserves not only spirals, but also position, normal vector, and curvature at each apex x of $\partial\mathcal{A}$. All leaves are centers of osculating disks at some apex x .

Let us remove from $\partial\mathcal{B}_n$ the containing circular arc b_x for each apex x whose osculating disk is included in \mathcal{B}_n (and hence is maximal for \mathcal{B}_n). This decomposes $\partial\mathcal{B}_n$ into components. The lengths of the arcs b_x shrink to zero as $\Omega(n^{-1})$ by construction of \mathcal{B}_n , as does their minimum d_n . Apart from disks for leaves, each maximal disk D_p for \mathcal{B}_n has contact to at least two different components. (Otherwise, there would be a supplementary leaf of $M(\mathcal{B}_n)$.) For such a disk D_p , we have the angle inequality $\xi_p \geq \xi_n$, for $\xi_n = 2 \arcsin(d_n/2L)$, and L denoting the geometric diameter of \mathcal{B}_n . Because $d_n \rightarrow 0$ as $\Omega(n^{-1})$ and since L is a constant, we have $1 - \cos(\xi_n/2) = \Omega(n^{-2})$. Moreover, $H(\partial\mathcal{A}, \partial\mathcal{B}_n) \rightarrow 0$ as $O(n^{-3})$ by construction. That is, the condition in Lemma 5 holds for almost all maximal disks D_p for \mathcal{B}_n when n is sufficiently large. Consequently, for each point

$p \in M(\mathcal{B}_n)$ there exists a point $q \in M(\mathcal{A})$ such that $\|p - q\| \rightarrow 0$ as $O(n^{-1})$. That is, the one-sided Hausdorff distance $\delta(M(\mathcal{B}_n), M(\mathcal{A}))$ converges at this speed.

The other direction can be proved similarly. For each apex x of \mathcal{A} , we define a neighborhood c_x on $\partial\mathcal{A}$ of length $n^{-3/4}$. Removal of the segments c_x leads us to two types of maximal disks D_q for \mathcal{A} , depending on whether D_q touches a single segment c_x (q is then close to x), or not. For the latter type, the analysis is the same as above, and shows that q approaches the center of some maximal disk for \mathcal{B}_n at speed $O(n^{-3/2})$. For the former type, due to Lemma 6, the distance $\|q - p\|$ between q and the leaf $p \in M(\mathcal{B}_n)$ associated with c_x behaves as $\Theta(n^{-3/4})^2$, i.e., the same. The one-sided Hausdorff distance $\delta(M(\mathcal{A}), M(\mathcal{B}_n))$ thus converges at that speed.

Note that the global convergence speed of the medial axis with respect to the Hausdorff distance is $\Theta(n^{-1})$, whereas the error of the boundary approximation improves as $\Theta(n^{-3})$. This is due to the behavior of the medial axis close to its leaves. When we restrict ourselves to the λ -medial axis [5] for *any* $\lambda > 0$, then d_n in the formula for ξ_n becomes a constant, and the approximation speed is $\Theta(n^{-3})$ by Lemma 5. This well compares to using a size- n point sample on $\partial\mathcal{A}$ and pruning its Voronoi diagram, as the approximation speed then is only $\Theta(n^{-1})$.

Acknowledgements. Thanks go to Raimund Seidel for discussions on Section 3.

References

1. Alt, H., Cheong, O., Vigneron, A.: The Voronoi diagram of curved objects. *Discrete & Computational Geometry* 34, 439–453 (2005)
2. Attali, D., Boissonnat, J.-D., Edelsbrunner, H.: Stability and computation of medial axes – a state-of-the-art report. In: Mller, T., Hamann, B., Russell, B. (eds.) *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, Springer Series on Mathematics and Visualization (to appear)
3. Avis, D., Toussaint, G.T.: An efficient algorithm for decomposing a polygon into star-shaped polygons. *Pattern Recognition* 13, 395–398 (1981)
4. Bhattacharya, B.K., El Gindy, H.: A new linear convex hull algorithm for simple polygons. *IEEE Trans. Information Theory* IT-30, 85–88 (1984)
5. Chazal, F., Lieutier, A.: Stability and homotopy of a subset of the medial axis. In: *Proc. 9th ACM Symp. Solid Modeling and Applications*, pp. 243–248 (2004)
6. Chazal, F., Soufflet, R.: Stability and finiteness properties of medial axis and skeleton. *J. Dynamical and Control Systems* 10, 149–170 (2004)
7. Chazelle, B.: A theorem on polygon cutting with applications. In: *Proc. 23rd IEEE Symp. FOCS*, pp. 339–349 (1982)
8. Chin, F., Snoeyink, J., Wang, C.A.: Finding the medial axis of a simple polygon in linear time. In: Staples, J., Katoh, N., Eades, P., Moffat, A. (eds.) *ISAAC 1995*. LNCS, vol. 1004, pp. 382–391. Springer, Heidelberg (1995)
9. Choi, H.I., Choi, S.W., Moon, H.P.: Mathematical theory of medial axis transform. *Pacific J. Mathematics* 181, 57–88 (1997)
10. Dobkin, D.P., Souvaine, D.L.: Computational geometry in a curved world. *Algorithmica* 5, 421–457 (1990)

11. Emiris, I.Z., Kakargias, A., Pion, S., Teillaud, M., Tsigaridas, E.P.: Towards an open curved kernel. In: Proc. 20th Ann. ACM Symp. Computational Geometry, pp. 438-446 (2004)
12. Farin, G.: Curves and Surfaces for Computer Aided Geometric Design. Academic Press, San Diego (1997)
13. Farin, G., Hoschek, J., Kim, M.-S.: Handbook of Computer Aided Geometric Design. Elsevier, Amsterdam (2002)
14. Garey, M.R., Johnson, D.S., Preparata, F.P., Tarjan, R.E.: Triangulating a simple polygon. Information Processing Letters 7, 175-179 (1978)
15. Graham, R.L.: An efficient algorithm for determining the convex hull of a finite planar set. Information Processing Letters 1, 132-133 (1972)
16. Graham, R.L., Yao, F.F.: Finding the convex hull of a simple polygon. J. Algorithms 4, 324-331 (1984)
17. Held, M., Eibl, J.: Biarc approximation of polygons with asymmetric tolerance bands. Computer-Aided Design 37, 357-371 (2005)
18. Hertel, S., Mehlhorn, K.: Fast triangulation of the plane with respect to simple polygons. Information & Control 64, 52-76 (1985)
19. Klein, R., Mehlhorn, K., Meiser, S.: Randomized incremental construction of abstract Voronoi diagrams. Computational Geometry: Theory and Applications 3, 157-184 (1993)
20. Kong, X., Everett, H., Toussaint, G.T.: The Graham scan triangulates simple polygons. Pattern Recognition Letters 11, 713-716 (1990)
21. Lee, D.T.: Medial axis transformation of a planar shape. IEEE Trans. Pattern Analysis and Machine Intelligence PAMI-4, 363-369 (1982)
22. Lee, D.T., Preparata, F.P.: Location of a point in a planar subdivision and its applications. SIAM J. Computing 6, 594-606 (1977)
23. McCallum, D., Avis, D.: A linear algorithm for finding the convex hull of a simple polygon. Information Processing Letters 9, 201-206 (1979)
24. Meek, D.S., Walton, D.J.: Approximation of a planar cubic Bézier spiral by circular arcs. J. Computational and Applied Mathematics 75, 47-56 (1996)
25. Meek, D.S., Walton, D.J.: Spiral arc spline approximation to a planar spiral. J. Computational and Applied Mathematics 107, 21-30 (1999)
26. Melkman, A.: On-line construction of the convex hull of a simple polygon. Information Processing Letters 25, 11-12 (1987)
27. Ong, C.J., Wong, Y.S., Loh, H.T., Hong, X.G.: An optimization approach for biarc curve fitting of B-spline curves. Computer-Aided Design 28, 951-959 (1996)
28. Ramamurthy, R., Farouki, R.T.: Voronoi diagram and medial axis algorithm for planar domains with curved boundaries I. Theoretical foundations. J. Computational and Applied Mathematics 102, 119-141 (1999)
29. Reif, U.: Uniform B-spline approximation in Sobolev spaces. Numerical Algorithms 15, 1-14 (1997)
30. Sabin, M.A.: The use of circular arcs to form curves interpolated through empirical data points. Rep. VTO/MS/164, British Aircraft Corporation (1976)
31. Šír, Z., Feichtinger, R., Jüttler, B.: Approximating curves and their offsets using biarcs and Pythagorean hodograph quintics. Computer-Aided Design 38, 608-618 (2006)
32. Yang, X.: Efficient circular arc interpolation based on active tolerance control. Computer-Aided Design 34, 1037-1046 (2002)
33. Yap, C.K.: An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments. Discrete & Computational Geometry 2, 365-393 (1987)