Algorithmica (1992) 8:365–389

Algorithmica

Simultaneous Inner and Outer Approximation of Shapes[1]

Rudolf Fleischer,[2] Kurt Mehlhorn,[2] Günter Rote,[3] Emo Welzl,[4]
and Chee Yap[5]

Abstract. For compact Euclidean bodies P, Q, we define $\lambda(P, Q)$ to be the smallest ratio r/s where $r > 0$, $s > 0$ satisfy $sQ' \subseteq P \subseteq rQ''$. Here sQ denotes a scaling of Q by the factor s, and Q', Q'' are some translates of Q. This function λ gives us a new distance function between bodies which, unlike previously studied measures, is invariant under affine transformations. If homothetic bodies are identified, the logarithm of this function is a metric. (Two bodies are *homothetic* if one can be obtained from the other by scaling and translation.)

For integer $k \geq 3$, define $\lambda(k)$ to be the minimum value such that for each convex polygon P there exists a convex k-gon Q with $\lambda(P, Q) \leq \lambda(k)$. Among other results, we prove that $2.118\ldots \leq \lambda(3) \leq 2.25$ and $\lambda(k) = 1 + \Theta(k^{-2})$. We give an $O(n^2 \log^2 n)$-time algorithm which, for any input convex n-gon P, finds a triangle T that minimizes $\lambda(T, P)$ among triangles. However, in linear time we can find a triangle t with $\lambda(t, P) \leq 2.25$.

Our study is motivated by the attempt to reduce the complexity of the polygon containment problem, and also the motion-planning problem. In each case we describe algorithms which run faster when certain implicit *slackness* parameters of the input are bounded away from 1. These algorithms illustrate a new algorithmic paradigm in computational geometry for coping with complexity.

Key Words. Polygonal approximation, Algorithmic paradigms, Shape approximation, Computational geometry, Implicit complexity parameters, Banach–Mazur metric.

1. Introduction. Most motion-planning problems, except for the simplest examples, have at least a quadratic time complexity in the worst case (see, for example, [14]). Our basic goal is to circumvent this apparent bottleneck by using heuristics. Yap [14] describes two general heuristics: the so-called *simplification heuristic* in which we try to replace a complicated robot body P by a simpler shape Q, and the *local expert heuristic* in which we invoke some specialized algorithm when the

---

[1] Work of all authors was partially supported by the ESPRIT II Basic Research Actions Program of the EC under Contract No. 3075 (project ALCOM). Rudolf Fleischer and Kurt Mehlhorn acknowledge also DFG (Grant SPP Me 620/6). Chee Yap acknowledges also DFG (Grant Be 142/46-1) and NSF (Grants DCR-84-01898 and CCR-87-03458). This research was performed when Günter Rote and Chee Yap were at the Freie Universität Berlin.
[2] Max-Planck-Institut Informatik (MPI), Im Stadtwald, W-6600 Saarbrücken, Federal Republic of Germany.
[3] Institut für Mathematik, Technische Universität Graz, Steyrergasse 30, A-8010 Graz, Austria.
[4] Institut für Informatik, Fachbereich Mathematik, Freie Universität Berlin, Arnimallee 2-6, 1000 Berlin 33, Federal Republic of Germany.
[5] Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012, USA.

Received June 3, 1990; revised May 15, 1991. Communicated by Mikhail J. Atallah.

robot is in some stereotyped environment (such as the vicinity of a door). Of course, the real challenge for theoretical robotics is to quantify precisely such heuristics. This paper provides a method for quantifying the simplification heuristic.

In real life we can often see instantly when a motion is possible or when a motion is impossible. This suggests that it may be possible to develop algorithms whose complexity reflects this phenomenon: it should run quickly for inputs where the possibility of a motion is "easy to see." Before we proceed to explain this idea, we should say that this idea is related to the concept of output-sensitive algorithms, but only in the sense that our algorithm also depends on some implicit complexity parameter of the input. After all, there does not seem to be an obvious connection between "easy-to-see-ness" and output size in our setting.

Let us formalize the idea of an implicit parameter. Assume that we want to move a convex polygon P amidst obstacles E from placement Z to Z'. We define the *slackness parameter* $s(P)$ $(=s(P, E, Z, Z'))$ to be the supremum of $s > 0$ such that there exists a motion for the body sP. Here sP denotes the scaling of P by s. To make this notion well defined, we assume that, in the initial and final positions Z and Z', P is surrounded by enough free space that the existence of a motion for sP is independent of the center of scaling for sP (as long as this center lies within P); see also Alt *et al.* [2]. Intuitively, we think of P moving from a large room to another large room through narrow doors and hallways.

Clearly, there exists a motion for P if and only if $s(P) \geq 1$. Now it is intuitively obvious that it is "easily seen" that no motion exists if the slackness parameter is very small (i.e., close to zero); likewise, it is "easily seen" that a motion exists if the slackness parameter is very large (i.e., $s(P) \gg 1$). When $s(P) \approx 1$, it is difficult to decide immediately whether a motion for P is possible.

What we would like to have is a simple substitute Q for P which should come as close as possible to satisfying the following conditions:

(i) If there is a motion for Q, then there is also a motion for P.
(ii) If there is no motion for Q, then there is no motion for P.

Of course, the only way to ensure these two conditions in general is to set $Q = P$. So we relax (ii):

(ii′) If there is no motion for Q, then there is no motion for P, *except* when it is "difficult to see" that there is a motion for P.

In other words:

(ii′) If there is no motion for Q, then $s(P) < 1$ or $s(P) \approx 1$.

We can make this more precise by choosing a constant $s_0 > 1$ and defining

$$s(P) \approx 1 \quad \Leftrightarrow \quad \frac{1}{s_0} < s(P) < s_0,$$

and thus our condition now becomes

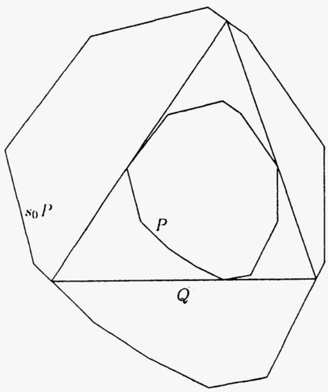

Fig. 1. A triangle Q is enclosed between P and s_0P.

(ii′) If there is no motion for Q, then $s(P) < s_0$ (i.e., there is no motion for s_0P).

We can achieve (i) by having P contained in Q, and we can achieve (ii′) by having Q contained in (a translate of) s_0P (see Figure 1).

Of course we can also weaken (i) instead of (ii), and we get:

(i′) If there is a motion for $\bar{Q}$, then $s(P) > 1/s_0$.
(ii) If there is no motion for $\bar{Q}$, then there is no motion for P.

We can achieve this by having $(1/s_0)P \subseteq \bar{Q}$ and $\bar{Q} \subseteq P$.

So if Q fulfills (i) and (ii′), then $\bar{Q} = (1/s_0)Q$ fulfills conditions (i′) and (ii), and vice versa. We see that both pairs of conditions lead in fact to the same approximation problem:

SHAPE APPROXIMATION PROBLEM. Given a convex figure P in the plane, find a "simple" polygon Q such that $P \subseteq Q \subseteq s_0P'$, where P' is a translate of P.

"Simple polygon" might for example mean triangle, quadrangle, rectangle, or ellipse. A set P and a scaled and translated copy $s(P + a)$ of it are called (*positively*) *homothetic*, for $s > 0$. It is clear that the roles of P and Q are interchangeable in the above statement, and thus we might as well look for a pair of homothetic simple polygons Q and s_0Q' for which $Q \subseteq P \subseteq s_0Q'$ holds (see Figure 2). This is in fact the formulation that we are going to work with. The two figures Q and s_0Q' approximate P from inside and from outside, motivating the title of our paper.

One of our results is that, for $s_0 = \frac{9}{4}$, we can find a *triangle* Q which fulfills the above relation in time linear in the number of vertices of the approximated polygon. For our application, this means that:

1. If there is no motion for the triangle Q, then there is no motion for P.
2. If there is a motion for s_0Q', then there is a motion for P.
3. Otherwise, the solution is not "easy to see" (i.e., $1/s_0 < s(P) < s_0$).

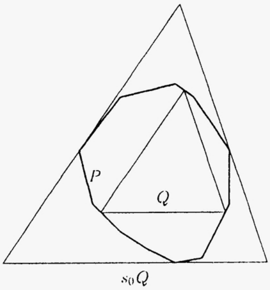

Fig. 2. A pair of triangles Q and s_0Q approximating a polygon P.

We can decide which of these three cases holds by running any known motion-planning algorithm for Q and for s_0Q'. Note that in the third case, we can continue testing with a good approximating k-gon, for $k = 4, 5, \ldots$. Thus we obtain an algorithm for motion planning whose running time degrades gracefully as the slackness parameter approaches 1. Note that we could run a standard motion-planning algorithm for P at the same time, so that the worst-case complexity can be guaranteed to be no more than the usual bound. (Alternatively, this can be guaranteed by letting k grow sufficiently fast.) This paradigm of making use of the slackness parameter depends on the ability to find good approximating pairs of k-gons efficiently.

By insisting that the two approximating polygons are homothetic we make sure that our application works even if rotations are not allowed in the robot motion. If we allowed rotations as well, then we could take better (i.e., smaller) values of s_0. For example, in the case of triangular approximations we could take $s_0 = 2$: Let t with vertices u, v, and w be a largest area triangle in the polygon P, and let T be the triangle with an edge through u parallel to vw, an edge through v parallel to uw, and an edge through w parallel to uv. Then $t \subseteq P \subseteq T$, and $-2t$ is a translate of T.

The above paradigm can be applied to the problem of polygon containment, again yielding an algorithm whose performance degrades gracefully as the implicit slack parameter approaches 1. Recall that the current fastest algorithm for placing a convex n-gon Q inside a convex m-gon P runs in time $O(nm^2)$ [3].

The preceding discussion motivates the following concept of approximation: For any two compact subsets Q, P of Euclidean space, let $\lambda(Q, P)$ be the infimum of the ratio r/s where $r, s > 0$ satisfy

$$sQ' \subseteq P \subseteq rQ''$$

and Q', Q'' are some translates of Q.

With this notion, our shape approximation problem can be formulated as follows:

> Given a convex figure P in the plane, find a "simple" polygon Q which minimizes $\lambda(Q, P)$.

In this paper we take "simple polygon" to mean a k-gon, for any fixed integer $k \geq 3$, and we study the maximum value of $\lambda(Q, P)$ which we may expect in the worst case.

In the Euclidean plane, for any integer $k \geq 3$, we define $\lambda(k, P)$ to be the infimum of $\lambda(Q, P)$ as Q ranges over all convex k-gons, and define $\lambda(k)$ to be the supremum of $\lambda(k, P)$ over all convex P. As mentioned above, we shall show that $\lambda(3) \leq \frac{9}{4}$.

For the distance measure $\lambda(P, Q)$, the size and position of P and Q are irrelevant; $\lambda(P, Q)$ depends only on the *shape* of P and Q. We have $\lambda(P, Q) = 0$ if and only if P and Q are positively homothetic. If homothetic bodies are identified, the logarithm of the function $\lambda(Q, P)$ turns out to be a metric, which is invariant under affine transformations. We study properties of this metric in Section 2.

Section 3 presents bounds for $\lambda(3)$, i.e., for approximation by triangles, and develops an $O(n^2 \log^2 n)$ algorithm for finding the best triangle approximation. In Section 4 we study the asymptotic behavior of $\lambda(k)$, and we discuss a number of open questions in Section 5. Finally, in an appendix we prove that the ratio between the area of a convex hexagon and the area of its largest contained triangle is at most 9/4; we need this result for our estimate of $\lambda(3)$.

This paper is an extended version of the conference paper [6].

2. A Metric on Shapes.

There are many different distance measures between convex bodies, like the Hausdorff distance, symmetric difference metric, perimeter deviation metric (see, for example, the survey by Gruber [8]). Typically, the definition of these metrics is motivated by the desire to approximate a convex body P in some Euclidean space by another body Q, where the metric function $d(P, Q)$ measures the quality of the approximation. The function $\lambda(P, Q)$ that we have defined in the introduction is different from the classical metrics in some important aspects. One notable property is that it is invariant under affine transformations: For any affine transformation τ, we have $\lambda(\tau Q, \tau P) = \lambda(Q, P)$. The classical metrics $d(Q, P)$ are invariant under rotations and translations, but not under other affine transformations. For example, if τ is a scaling by the factor μ, then we usually have

$$d(\tau Q, \tau P) = |\mu|^q \cdot d(Q, P),$$

where q is some integer between 1 and the dimension.

With suitable precautions, the logarithm of λ is a metric. By a (planar) *body* P we mean a compact subset of the plane. For any body P, let $\tilde{P}$ denote the class of bodies equivalent to P under translation and positive scaling, i.e., the homothets of P. We call such equivalence classes *shapes*. We say two bodies P, Q have the *same shape* if $\tilde{P} = \tilde{Q}$.

We first observe that λ is in fact a function on shapes: that is, if P, P' have the same shape and if Q, Q' have the same shape, then $\lambda(P, Q) = \lambda(P', Q')$. Hence the notation $\lambda(\tilde{P}, \tilde{Q})$ is meaningful. The following theorems are easy to prove.

THEOREM 2.1. *The function* $\tilde{\lambda}(\tilde{P}, \tilde{Q}) := \log \lambda(\tilde{P}, \tilde{Q})$ *defines a metric on shapes.*

THEOREM 2.2. *The functions λ, $\tilde{\lambda}$ are invariant under affine transformations, that is, if τ is an affine transformation, then $\lambda(P, Q) = \lambda(\tau P, \tau Q)$.*

The metric $\tilde{\lambda}$ has also been used by Kannan *et al.* [10] under the name of the Banach–Mazur metric. (It is a variation of the classical Banach–Mazur distance which applies to centrally symmetric bodies and allows arbitrary affine transformations, not just scalings.)

3. Approximation by Triangles.

In this section we give lower and upper bounds for $\lambda(3)$ and present an algorithm which constructs an optimal triangle approximation for a given n-gon in time $O(n^2 \log^2 n)$.

First, it is useful to introduce the following general concept: for two bodies Q, Q' with the same shape, and for a body P, we call (Q, Q') an *approximating pair* for P if $Q \subseteq P \subseteq Q'$; the *expansion factor* of (Q, Q') is the factor by which Q must be scaled in order for it to have the same size as Q'. Given a body P and a triangle t contained in P, let $T_P(t)$ be the smallest triangle of the same shape as t that contains P. Then $(t, T_P(t))$ is an approximating pair for P, and its expansion factor is denoted by $\chi_P(t)$. Note that $\lambda(t, P) \leq \chi_P(t)$, where equality holds if and only if t is a largest triangle of its shape contained in P.

The Maximum Area Heuristic. Computing the value $\lambda(k, P)$ (for k and P) seems to amount to finding an approximating pair (Q, Q') for p such that Q is a k-gon and the expansion factor equals $\lambda(k, P)$. There is a natural candidate for (Q, Q'), namely where Q is the largest convex k-gon contained in P. The next theorem shows how well this *maximum area heuristic* performs in case of the triangle.

THEOREM 3.1. *For any convex body C, any largest triangle t contained in C has the property that*

$$\lambda(t, C) \leqslant 9/4.$$

PROOF. We can apply an affine transformation that maps t to an equilateral triangle with unit side length; so we may assume that t is equilateral with unit side length from the beginning. Each edge of the triangle $T = T_C(t)$ touches C in at least one point; this gives three points which together with the vertices of t form a polygon P with at most six vertices; in general this will be a hexagon (see Figure 3). Let h and H denote the heights of t and T, respectively, and let $d := d_1 + d_2 + d_3$ where the d_i are the distances between corresponding edges of t and T. Then $H = h + d$, because t is equilateral. Let us denote by area$Q$ the area of body Q. We have area$t = h/2$ and area$P$ = area$t + \frac{1}{2}(d_1 + d_2 + d_3) = H/2$. So the expansion factor of (t, T) equals H/h = area$P$/area$t$. In the appendix we show that area$P$/area$t \leq 9/4$ if P is a convex polygon with at most six vertices, and t is a triangle of largest area contained in P. The theorem follows. □

COROLLARY 3.2. *For any convex n-gon, we can find in $O(n)$ time a triangular approximating pair (t, T) with expansion factor at most 9/4.*

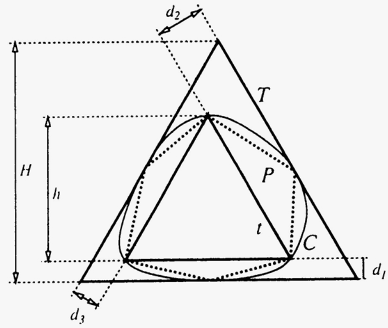

Fig. 3. The hexagon P in the proof of Theorem 3.1.

This follows from the fact that the largest area triangle contained in a convex polygon can be found in linear time [4]. We will see below that the maximum area heuristic is, in general, suboptimal.

Rigid Approximating Pairs. For the construction of an optimal triangle approximation of a given convex n-gon P, we look for a triangle t contained in P that minimizes $\chi_P(t)$; clearly, this triangle also minimizes $\lambda(t, P)$ and thus determines $\lambda(3, P)$. In a first step we reduce the set of possible candidates for t.

Let (t, T) be an approximating pair of triangles for P. A pair (v, V) of corresponding vertices of (t, T) is *free* if (i) $v \notin \partial P$ or (ii) $v \in \partial P$, but it is not a vertex of P and the two edges of T incident to V are not flush with edges of P (∂P denotes the boundary of P). We call (t, T) *rigid* if it has no free vertex pair and $T = T_P(t)$.

For example, in the pair of triangles in Figure 2, every pair of corresponding vertices is free. In Figure 4 (u, U) and (w, W) are not free for the pair $(\triangle uvw, \Delta UVW)$, but (v, V) is free. The following two lemmata will lead to a theorem that justifies that we restrict our attention to rigid approximating pairs.

LEMMA 3.3. *Let t be a triangle contained in P such that $\chi_p(t) = \lambda(3, P)$. Then all vertices of t lie on the boundary ∂P of P.*

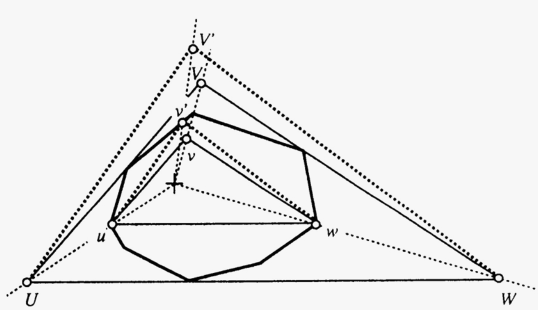

Fig. 4. Illustrating the proof of Lemma 3.3.

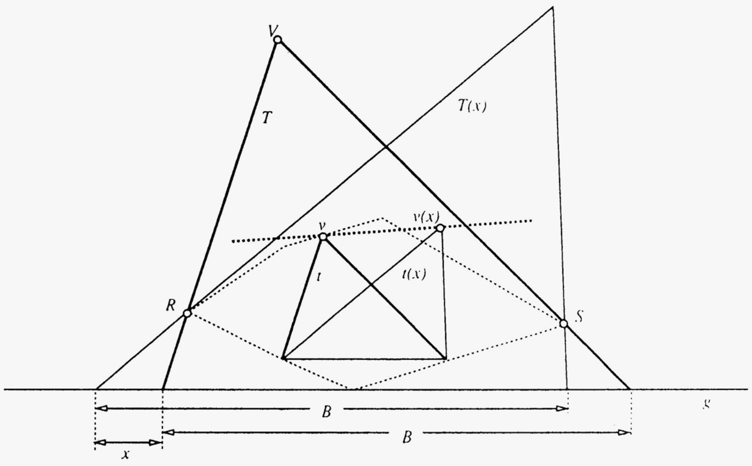

Fig. 5. Illustrating Lemma 3.4 (and the proof of Theorem 3.5).

PROOF. See Figure 4. Suppose that $t = \triangle uvw$ has vertex v not in ∂P. Consider $T = T_P(t)$. Without loss of generality, we assume that $T = \lambda t$, i.e., the scaling center of t and T is the origin. Since t is the largest triangle of its shape in P, μt is not in P for any $\mu > 1$; therefore, either μu or μw lies outside of P (not on ∂P) for every $\mu > 1$, say this holds for μu. Hence, the vertex $U = \lambda u$ of T corresponding to u has to lie outside of P.

Consider now a point v' on the boundary of P such that the triangle $t' = \Delta uv'w$ contains t and has v in its interior. Let $T' = \Delta UV'W$, where $V' = \lambda v'$. Then (t', T') is an approximating pair for P with the same expansion factor as (t, T). However, since the edge UV' does not touch P, it cannot be optimal; we use here that $U \notin \partial P$, and that the rest of edge UV' does not even touch T, since V lies in the interior of T'. □

The following lemma can be proved "directly" by some analytic calculations; we present a short proof (using cross-ratios) based on notes by Rolfdieter Frank [7]. It turns out that this is essentially a new *Schließungssatz* equivalent to Pappus' theorem as was pointed out by Armin Saam [12]. The lemma is illustrated in Figure 5, in a way that indicates already how we want to use it.

LEMMA 3.4. *Let T be a triangle with a base of length B, let R and S be points on the other two edges of T (but not on the base), and let g be the line which contains the base of T. If we move the base of T on g to the left (or right) by an amount x while preserving its length, this new base together with the (fixed) points R and S defines a triangle $T(x)$. Then, in a corresponding similar triangle $t(x)$ with a fixed base, the third vertex $v(x)$ (i.e., the vertex not on the base) moves on a straight line as x varies.*

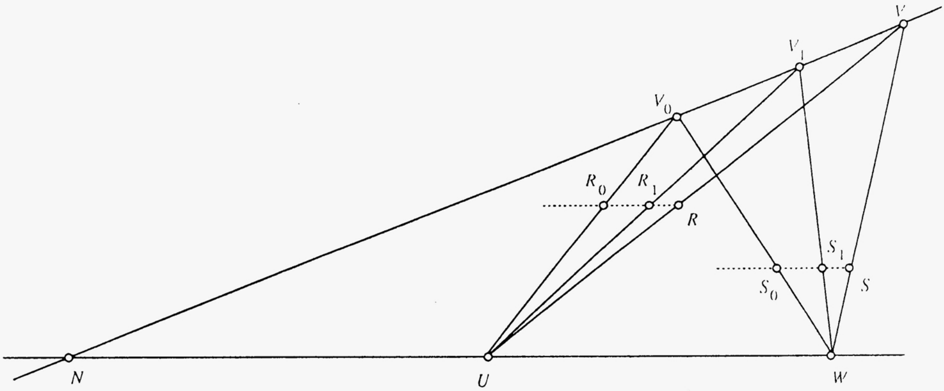

Fig. 6. Illustrating the proof of Lemma 3.4.

PROOF. Instead of moving the base of T and keeping the points R and S fixed, we keep the basis fixed and move the points R and S horizontally by equal amounts, see Figure 6. We want to show that V moves on a line. We proceed as follows: assume that the base UW is horizontal and take two triangles $\triangle UV_0W$ and $\triangle UV_1W$ "generated" by points R_0, S_0 and horizontally shifted points R_1, S_1, respectively. Now we move the point V on the line through V_0, V_1 and watch the intersections R and S of the sides UV and VW with the horizontal lines through R_0R_1 and S_0S_1, respectively. We will prove the lemma by showing that the distances R_0R and S_0S are equal.

Let us recall from projective geometry that the *cross-ratio* (A, B, C, D) of four collinear points is defined as

$$(A, B, C, D) = \frac{\overrightarrow{AC}}{\overrightarrow{AD}} \cdot \frac{\overrightarrow{BD}}{\overrightarrow{BC}},$$

where the sign of the directed distances $\overrightarrow{AC}$, $\overrightarrow{AD}$, $\overrightarrow{BC}$, and $\overrightarrow{BD}$ has to be taken th respect to a fixed chosen orientation of the line containing the points; the value of the cross-ratio is independent of this choice of orientation. If one of the points, for example A, is at infinity, the value of $\overrightarrow{AC}/\overrightarrow{AD}$ is taken as 1. In this case, the cross-ratio (∞, B, C, D) simplifies to the ratio $\overrightarrow{BD}:\overrightarrow{BC}$. It is an elementary fact of projective geometry that the cross-ratio is invariant under central projections.

To use this fact for our proof, we denote by N the intersection of the base line through UW with the line on which vertex V moves. If we project the four points N, V_0, V_1, and V on the horizontal lines through R_0R_1 and S_0S_1 from the centers U and W, respectively, we get

$$\overrightarrow{R_0R}:\overrightarrow{R_0R_1} = (\infty, R_0, R_1, R) = (N, V_0, V_1, V) = (\infty, S_0, S_1, S) = \overrightarrow{S_0S}:\overrightarrow{S_0S_1}.$$

However, $\overrightarrow{R_0R_1}$ and $\overrightarrow{S_0S_1}$ are equal by choice, and thus $\overrightarrow{R_0R}$ and $\overrightarrow{S_0S}$ are also equal. □

THEOREM 3.5. *For any convex polygon P, there is a rigid approximating pair (t, T) with expansion factor $\lambda(3, P)$.*

PROOF. Let (t, T) be an approximating pair with expansion factor $\lambda(3, P)$. We know from Lemma 3.3 that every vertex of t is on δP. Suppose there is a free vertex pair (v, V). Then each of the two edges of T incident to V touches P in one point only, say in point R and in point S, respectively. If R is a vertex of T, then we can make T smaller by rotating the edge RV about R until it is flush with an edge of P. We get a triangle T', which forms with the analogously manipulated triangle t' in P an approximating pair with the same expansion factor, and this approximating pair has one free vertex pair less. Continuing to transform the pair in this way we obtain a rigid pair.

So we may assume that R and S are not vertices of T, and the assumptions of Lemma 3.4 are satisfied. We move the base (the edge not containing V) on its line while preserving contact with P in points S and R. Until an edge gets flush with an edge of P, the vertex v of a similar triangle with the same base as t moves on a line. If this line intersects the interior of P, then we get a contradiction to the optimality of (t, T) via Lemma 3.3 (see Figure 5). Thus this line contains the edge on which v sits and we can perform this motion while preserving the expansion factor, until either v meets a vertex of P or one of the two edges of T containing R and S, respectively, gets flush with an edge of P; then (v, V) is not free. □

So we can find the optimal approximating pairs of triangles by considering only rigid approximating pairs. Every rigid approximating pair (t, T) falls into at least one of the following classes:

A. All three vertices of t are also vertices of P.
B. One vertex of t is also a vertex of P and the opposite edge of T is flush with an edge of P.
C. Two edges of T are flush with edges of P.

These classes are not disjoint but we will find the overall optimum by computing the optimum of each class. The pairs in each class can be further classified by the incidences between the edges and vertices of the inner and outer triangle on one side, and of the approximated polygon on the other side. We call this the *type* of a rigid approximating pair.

For Class A a type is specified by the three vertices of P which are the vertices of t; this completely determines (t, T), because $T = T_P(t)$. The situation gets slightly more subtle in the other classes, because the type does not always completely determine the triangles, and a type may contain a continuous family of solutions. Fortunately, we only have to deal with one-parametric families of solutions, and thus it is easy to find the optimum.

The type of a Class B pair (t, T) is specified by a vertex v of t and the vertex of P which coincides with v, the edge e of P which is contained in the edge of T opposite to V (V is the vertex of T corresponding to v) and the portions (edges or vertices) of ∂P which contain the other vertices of t, and the portions which

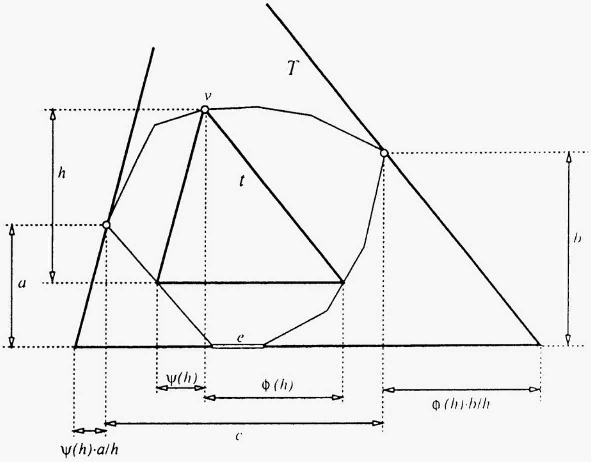

Fig. 7. How to compute the optimum of a given type in Class B.

are contained in the other edges of T. There is one case where the type does not uniquely determine the triangles t and T, namely when the other vertices (different from v) of t are both on edges of P (not vertices), and the other edges of T both touch P in vertices (not edges). However, the expansion factor of a pair (of fixed type) is a rational function of the form $(c_1h + c_2)/(c_3h^2 + c_4h)$ in the height h of the smaller triangle (the c_i's are constants depending on the type); this should be clear from Figure 7. Note that ϕ and ψ are linear functions in h. The base of t has length $\psi(h) + \phi(h)$, and the base of T has length $\psi(h)a/h + \phi(h)b/h + c$, where a, b, and c are constants. The expansion factor is the ratio of these two expressions. Hence the optimum of a given type can be found by solving a quadratic equation and checking whether the roots of this equation correspond to valid pairs (t, T) of the type. It may happen that the optimum occurs at a boundary position of the type, but this would already be a different type, and therefore such a solution is taken care of correctly.

In Class C we have two edges e and f of P which are contained in edges UV and VW of T. This also determines the vertex V. In addition, to completely specify the type, we have to say on which vertex or edge of P each vertex of t lies, and which vertex or edge of P is touched by the edge UW. The only case in which the solution is not unique occurs when the vertices u, v, and w of t are contained in edges (not vertices) of P, and the edge UW of T touches P at some vertex R (not at an edge). In this case the distance $\overline{VR}$ is constant (see Figure 8). Therefore, $\chi_P(t)$ is inversely proportional to the corresponding distance vr in the small triangle t. Let us assume that the line VR is vertical, and let d be the distance of v from the next vertex of P on the left. The coordinates (u_x, u_y), (v_x, v_y), and (w_x, w_y) of the points u, v, and w are then linear functions of d. By dividing the triangle t into the triangles $\triangle vrw$ and $\triangle vru$ we can see that the area of t is half the length of vr

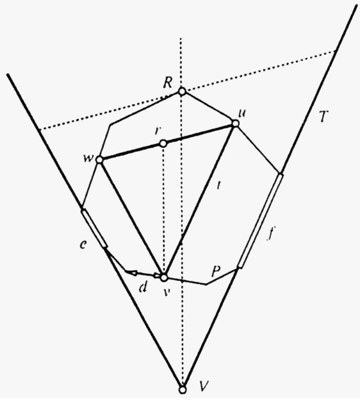

Fig. 8. How to compute the optimum of a given type in Class C.

times the horizontal distance between w and u. This gives the following formula for $\overline{vr}$:

$$\overline{vr} = \frac{\begin{vmatrix} u_x - v_x & w_x - v_x \\ u_y - v_y & w_y - v_y \end{vmatrix}}{u_x - w_x}.$$

Setting the derivative of this rational function to zero yields a quadratic equation for d, and thus the minimization problem for this type can be solved by elementary means.

An Operation on Triangles. For the analysis of the algorithm we are going to describe, we need the following argument. If (t_0, T_0) and (t_1, T_1) are approximating pairs of a convex polygon P with the same expansion factors, then—under certain circumstances to be specified later—we can continuously deform one pair into the other via approximating pairs (t_μ, T_μ), $0 \le \mu \le 1$, while preserving the same expansion factor. For $t_0 = \triangle u_0 v_0 w_0$ and $t_1 = \triangle u_1 v_1 w_1$, we define t_μ as the triangle $\triangle((1-\mu)u_0 + \mu u_1)((1-\mu)v_0 + \mu v_1)((1-\mu)w_0 + \mu w_1)$ and we define T_μ in an analogous way as an intermediate triangle between T_0 and T_1; note that in the definition of t_μ the order of the vertices defining t_0 and t_1 makes a difference for the resulting t_μ.

Observe right away that t_μ and T_μ have the same shape, that (t_μ, T_μ) has the same expansion factor as (t_0, T_0) and (t_1, T_1), and that t_μ is contained in P. The first two properties can be easily checked by plugging in the formulas, and the last property follows from the fact that the vertices of t_μ are contained in the convex hull of $t_0 \cup t_1$.

However, in general, T_μ will not contain P. The following lemma, though, will ensure that in the situations we consider.

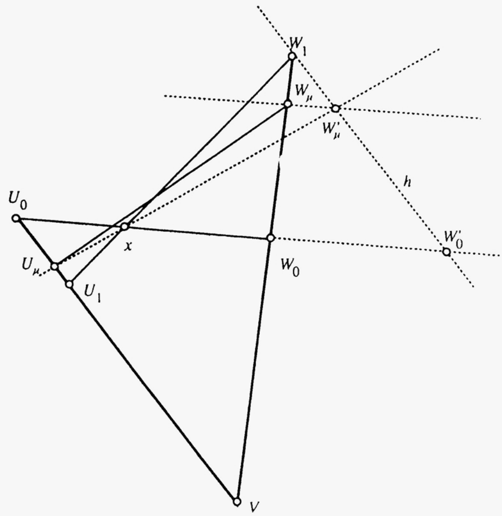

Fig. 9. Illustrating Lemma 3.6 and its proof.

LEMMA 3.6 (see Figure 9). *Let V, U_0, W_1 be three noncollinear points in the plane, let U_1 be a point on the open segment VU_0, and let W_0 be a point on the open segment VW_1. For μ, $0 < \mu < 1$, we define $U_\mu = (1-\mu)U_0 + \mu U_1$ and $W_\mu = (1-\mu)W_0 + \mu W_1$. Then $\triangle U_\mu V W_\mu$ contains $\triangle U_0 V W_0 \cap \triangle U_1 V W_1$ for all $0 \le \mu \le 1$.*

PROOF. It suffices to prove that $\triangle U_\mu V W_\mu$ contains the point x of intersection between segments $U_0 W_0$ and $U_1 W_1$. Consider the line h parallel to $U_0 V$ containing W_1, let W'_0 be the point where h intersects the line containing $U_0 W_0$, and let $W'_\mu = (1-\mu)W'_0 + \mu W_1$. Then $U_\mu W'_\mu$ contains x. However, W'_μ can be obtained also by intersecting the line parallel to $W_0 W'_0$ through W_μ with h. The lemma now follows easily. □

COROLLARY 3.7. *If (t_0, T_0) and (t_1, T_1) are approximating pairs of P such that T_0 has two edges flush with edges of P and T_1 has two edges flush with the same two edges of P, then all (t_μ, T_μ) for $0 \le \mu \le 1$ are approximating pairs of P.*

COROLLARY 3.8 (see Figure 10). *Let $T_0 = \triangle U_0 V_0 W_0$ and $T_1 = \triangle U_1 V_1 W_1$ be two triangles such that the base segment $U_0 W_0$ of T_0 is contained in the base segment $U_1 W_1$ of T_1 and the vertex V_1 is contained in T_0. Then T_μ contains $T_0 \cup T_1$ for all $0 \le \mu \le 1$.*

PROOF. The line through V_0 and V_1 intersects the basis $U_0 V_0$ and thus separates each of the triangles T_0 and T_1 into two triangles, to which Lemma 3.6 can be applied. □

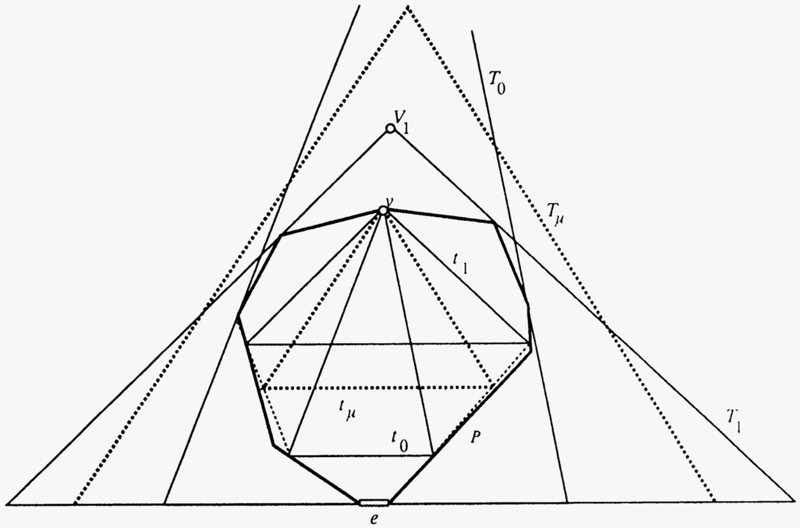

Fig. 10. Illustrating the proof of the unimodularity of $\chi(h)$ in Class B.

Computing an Optimal Approximating Pair. The algorithms for each class are as follows:

For Class A we investigate all pairs u, w of vertices of P and look for a third vertex v which minimizes $\chi_P(\triangle uvw)$. For fixed vertices u and w, we abuse notation and write $\chi(x)$ as short for $\chi_P(\triangle uxw)$, where x is a point in P. We want to show that $\chi(x)$ is unimodal as x moves from u to w on ∂P (on either side of uw). Suppose $\chi(x') = \chi(x'')$ for two points x' and x'' on ∂P, both on the same side of uw. Let $T' = \Delta U'X'W'$ be the triangle $T_P(\triangle ux'w)$, and let $T'' = \triangle U''X''W'' = T_P(\triangle ux''w)$.

Now we apply Lemma 3.4. The bases of T' and T'' lie on the same line, and since $(\triangle ux'w, T')$ and $(\triangle ux''w, T'')$ have the same expansion factors, the bases of T' and T'' have the same lengths. Since all edges of T' and T'' touch P, the segment $U'X'$ must intersect $U''X''$ in a point R, and $W'X'$ must intersect $W''X''$ in a point S. Now move the base of T' toward the base of T'', and consider the triangles with this base and the other two edges containing R and S; all these triangles contain $T' \cap T'' \supseteq P$. As we observe similar copies with base uw, Lemma 3.4 tells us that the third vertex moves monotonically on a line from x' to x''. Hence every point x on ∂P between x' and x'' satisfies $\chi(x) \leq \chi(x')$, and so χ is unimodal.

LEMMA 3.9. *The optimal approximating pair in Class A can be computed in time* $O(n^2 \log^2 n)$.

PROOF. For every pair u, w of vertices P, and for both sides of uw, we perform a Fibonacci search for the vertex v that minimizes $\chi_P(\triangle uvw)$. This needs $O(\log n)$ vertices to be visited, and for each such vertex v we have to compute $T_P(\triangle uvw)$, which takes $O(\log n)$ time. This gives the claimed time bound. □

In Class B we fix an edge e and a vertex v of P, and consider all triangles $\triangle uvw$, with u, w on ∂P, and uw parallel to e. We write $\chi(h)$ as short for $\chi_P(\triangle uvw)$, if h is the height of $\triangle uvw$ at base uw. Again we show that $\chi(h)$ is unimodal.

Let $h_0 > h_1$ with $\chi(h_0) = \chi(h_1)$. For $i = 0, 1$, let t_i be the triangle $\triangle u_i v w_i$ with height h_i, with u_i, w_i on ∂P, and $u_i w_i$ parallel to e, and let $T_i = T_P(t_i)$. Note that the base of T_0 (containing e) is completely contained in the base of T_1, and that the third vertex V_1 of T_1 is contained in T_0 (see Figure 10). Hence, the line through V_1 and the third vertex V_0 of T_0 intersects the bases of T_0 and T_1. By Corollary 3.8 we know that the triangles T_μ contain $T_0 \cap T_1 \supseteq P$ for all μ, $0 < \mu < 1$. Now it easily follows that $\chi(h) \leq \chi(h_0)$ for all h between h_0 and h_1.

Lemma 3.10. *The optimal approximating pair in Class B can be computed in time* $O(n^2 \log^2 n)$.

Proof. For every pair of an edge e and a vertex x of P, we search for the type that contains the optimal triangle $t = \triangle uvw$ with one vertex v on x, and the opposite edge uw parallel to e. First we perform a Fibonacci search for the optimal position for the vertex u of T among the vertices of P. This will limit the possible positions of u to one vertex and two edges of P. Then we similarly identify two possible edges and a possible vertex for w. Finally, we search among the edges of P where the edge UV of T may lie flush, again using Fibonacci search. This gives us a possible edge and two possible vertices where this edge touches. After a similar search for the edge VW, we find the optimum in each of the resulting types in constant time, either by computing the ratio for the unique solution or by optimizing locally as described above (see Figure 7). □

Finally, we have reached Class C. Let e and f be edges of P. For point v on ∂P, consider the triangle $t = \triangle uvw$ with u, w on ∂P, uv parallel to e, and vw parallel to f. We write $\chi(v)$ as short for $\chi_P(t)$. Using Corollary 3.7 it is now easy to show that $\chi(v)$ is unimodal as v moves on ∂P between e and f in the part where $T_P(t)$ has edges flush with e and f.

Lemma 3.11. *The optimal approximating pair in Class C can be computed in time* $O(n^2 \log^2 n)$.

Proof. For every pair e, f of edges of P we select a few possible types of approximating pairs where the outer triangle has edges flush with e and f. This is carried out similarly as in Class B: Let V be the common vertex of these two outer edges, as in Figure 8. By first searching among all vertices of P as possible positions for v we identify the two possible edges and one possible vertex where v may lie. This is done by Fibonacci search in $O(\log^2 n)$ time. Subsequently, we identify two possible edges and a possible vertex for u, and then for w; each by a separate Fibonacci search. Finally, we search among the edges of P where the third edge UW of T may lie flush. This gives us a possible edge and two possible vertices where this edge touches. We end up with a small number of types, which have either a unique solution, or can be handled as described before. □

We conclude with:

THEOREM 3.12. *Given an n-gon P we can compute an optimal triangular approximating pair* (t, T) *and the value* $\lambda(3, P)$ *in time* $O(n^2 \log^2 n)$.

Approximating the Regular Pentagon. We conclude this section by determining the optimal approximating pair for the regular pentagon. It will provide us with a lower bound for $\lambda(3)$. Somewhat surprisingly, this bound is tighter than the bound $\lambda(3, D) = 2$ for a disk D.

The optimal approximation (t, T) for a regular pentagon turns out to be in Class B, with a common vertex v of t and the pentagon, and an edge of T flush with the edge of the pentagon opposite to v. All other types and classes are boundary cases of this type or symmetric to it, or they can be dismissed as worse by direct calculations. Figure 11 depicts the optimal situation for a regular pentagon, where the indicated distances refer to a pentagon of side length 1. The distances are labeled according to Figure 7. The slope of the lower right edge of the pentagon is $1/\tan 18°$, the vertical distance of the base of t from the base of the pentagon is $H - h = (\cot 18°)/2 - h$. Thus $\psi(h) = \phi(h) = \frac{1}{2} + (H - h) \tan 18° = 1 - h \tan 18°$.

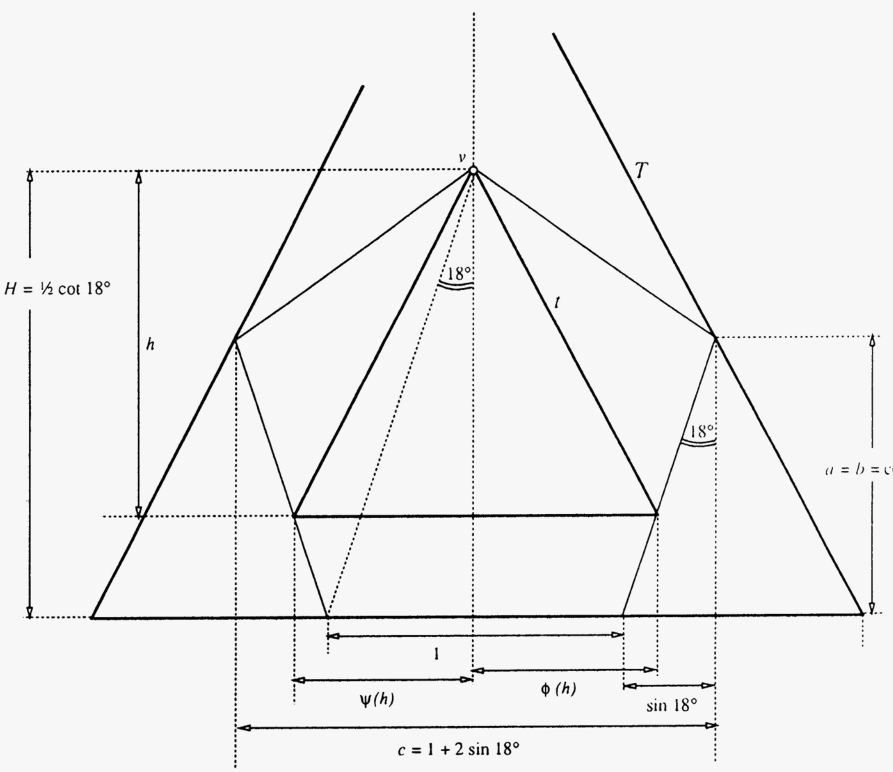

Fig. 11. The optimal approximation for the regular pentagon.

The formula derived above for Class B gives

$$\text{(1)} \qquad \frac{\psi(h)a/h + \phi(h)b/h + c}{\psi(h) + \phi(h)} = \frac{(2 - 2h \tan 18°) \cos 18°/h + 1 + 2 \sin 18°}{2 - 2h \tan 18°} = \frac{h + 2 \cos 18°}{2h - 2h^2 \tan 18°}.$$

By setting the derivative to zero and solving the resulting quadratic equation we get the optimal value of h:

$$h = \cos 18°\left(-2 + \sqrt{4 + \frac{2}{\sin 18°}}\right) = \cos 18°(\sqrt{6 + 2\sqrt{5}} - 2) = \cos 18°(\sqrt{5} - 1),$$

using the value of $\sin 18° = (\sqrt{5} - 1)/4$. Substituting this into expression (1) gives an expansion factor of $1 + \sqrt{5}/2$.

THEOREM 3.13.

$$2.118\ldots = 1 + \frac{\sqrt{5}}{2} \le \lambda(3) \le \frac{9}{4}.$$

We conjecture that the regular pentagon is indeed the worst case for approximation by triangles, and that $\lambda(3)$ is equal to $1 + \sqrt{5}/2$.

4. Upper Bounds for k-gons. The maximal area heuristic applied to quadrilaterals can be seen to yield an approximating pair (Q, Q') with an expansion factor at most 2. Hence $\lambda(4) \le 2$, but, in fact, Schwarzkopf *et al.* [13] have recently shown that we may assume that Q is rectangular and such a rectangular approximating pair can be computed in time $O(\log^3 n)$ if P is a polygon with its n vertices given sorted in a linear array. Furthermore, this bound of 2 is optimal when restricted to rectangular approximation.

The rest of this section is devoted to the asymptotic behavior of $\lambda(k)$. Any disk D can be approximated by a regular k-gon Q with $\lambda(Q, D) = 1/\cos(\pi/k)$, which is optimal. This gives us a lower bound on $\lambda(k)$:

LEMMA 4.1. $\lambda(k) \ge 1/\cos(\pi/k) > 1 + \pi^2/(2k^2)$.

In fact this lower bound on $\lambda(k) - 1$ is tight up to a constant factor, as will be shown below. The idea of the proof is to reduce our approximation problem to approximation with respect to the Hausdorff distance, for which an $O(1/k^2)$ bound is known. For any two bodies P and R, their Hausdorff distance $d(P, R)$ is defined as follows:

$$d(P, R) = \max\left\{\sup_{x \in P} \inf_{y \in R} \overline{xy}, \sup_{y \in R} \inf_{y \in P} \overline{xy}\right\}.$$

We invoke the following result:

LEMMA 4.2 (see [8]). *Let a convex body P of perimeter U be given, and let $k \geq 3$. Then P contains a k-gon R with*

$$d(P, R) \leq U \frac{\sin(\pi/k)}{2k} < U \frac{\pi}{2k^2}.$$

However, for polygons that are very long and thin, the Hausdorff distance d is not a good approximation to our distance λ. This is because d measures the actual Euclidean distance from each point to the nearest point in the other body, whereas λ measures a relative scaling factor so that the effect of a point on λ is somehow inversely proportional to its distance from the scaling origin. Therefore, first we have to apply an affine transformation to our body P to make it roughly "round." We expect that for "round" bodies, there will not be too much difference between the Hausdorff distance and our distance measure λ. The following lemma makes this precise.

LEMMA 4.3. *Let P be a convex body containing a convex body R such that $d(P, R) \leq \varepsilon$, and suppose that R contains a disk of radius a. Then $\lambda(P, R) \leq 1 + \varepsilon/a$.*

PROOF. Let us take the center of the disk of radius a as our origin O. We claim that $P \subseteq (1 + \varepsilon/a)R$, where the scaling of R is centered at O. To see this, look at a half-line h emanating from O which intersects the boundaries of R and P in points r and p, respectively (see Figure 12). Let us draw a supporting line of R through r, and denote the points on this line closest to O and p by O' and p', respectively. By considering similar triangles, we can conclude that

$$\frac{\overline{rp}}{\overline{rO}} = \frac{\overline{pp'}}{\overline{OO'}}.$$

However, $\overline{pp'} \leq \varepsilon$ and $\overline{OO'} \geq a$, by our assumptions, and thus we get

$$\frac{\overline{Op}}{\overline{Or}} = \frac{\overline{Or} + \overline{rp}}{\overline{Or}} \leq 1 + \frac{\varepsilon}{a}. \qquad \square$$

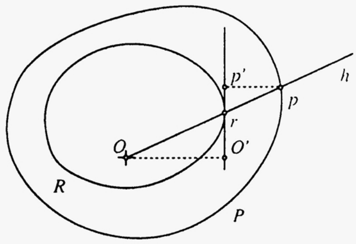

Fig. 12. Illustrating the proof of Lemma 4.3.

LEMMA 4.4. *For* $k \geq 5$, $\lambda(k) \leq 1 + 2\pi^2/(k^2 - 2\pi^2)$.

PROOF. Let P be a convex body to be approximated by a k-gon. We know by a result of John [9] (see also [11]) that there is an ellipse $E \subseteq P$, such that $P \subseteq E'$, where E' is the ellipse E scaled by a factor of 2 about its center. Now we choose an appropriate affine transformation such a that E becomes a unit disk. Since P is contained in E' and since E' has perimeter 4π we know that the perimeter of P is at most 4π. By Lemma 4.2, we can find a k-gon R contained in P such that $d(P, R) \leq 4\pi^2/(2k^2)$. Let $\varepsilon := 2\pi^2/k^2$. Since P contains the unit disk E, R must contain a disk of radius $1 - \varepsilon$, and applying Lemma 4.3 with $a = 1 - \varepsilon$ gives the result. (Note that $1 - \varepsilon > 0$ for $k \geq 5$.) □

We can summarize the results of this section as follows:

THEOREM 4.5. $\lambda(k) = 1 + \Theta(1/k^2)$.

5. Further Research.

- The algorithm for constructing an optimal approximating triangle is slow but it is not clear how the exhaustive search used in Theorem 3.12 can be made considerably faster.
- Even the triangle approximation of the regular pentagon (lower bound in Theorem 3.13) is not fully understood. It turns out that the scaling center for the optimal approximating triangle pair is just the midpoint of the pentagon but we do not know if there is a deeper reason for this.
- For small k we would like to have explicit tight bounds on $\lambda(k)$ instead of the asymptotic bounds in Section 4. Furthermore, we would like to find algorithms which efficiently construct optimal (or nearly optimal) k-gons. One candidate for such an algorithm is the maximum area heuristic. Currently we do not have general bounds on the performance of this heuristic.
- What can be said about the minimum enclosing polygon [1] heuristic? Again we know that in general it is not optimal (the example of the regular pentagon again).
- We have not considered higher-dimensional problems.

Appendix. Largest Area Triangles in Hexagons. We want to prove that for a convex polygon P with at most six vertices the ratio between its area and the area of the largest triangle contained in P is at most 9/4. For a convex body C, we denote the ratio between its area and the area of the largest triangle contained in it by ratio$C$. For arbitrary convex bodies C (without restriction on the number of vertices) a tight bound of ratio$C \leq 4\pi/(3\sqrt{3}) = 2.41\ldots$ is known [5].

We need some notation. Given a convex polygon P, a *critical triangle* is a triangle of largest area whose vertices are vertices of P. It is easily seen that there is always a critical triangle that has largest area among all triangles contained in

P. If P is a hexagon with vertices A, B', C, A', B, C' in clockwise order, then a triangle t is called *alternating* if $t = \triangle ABC$ or $t = \triangle A'B'C'$, and it is called *diagonal* if AA', BB', or CC' is an edge of t. Note that if t is neither alternating nor diagonal, then it is formed by three consecutive vertices of P.

A *largest-ratio instance* is a convex polygon P with at most six vertices that maximizes ratio$P$. A compactness argument shows that such a largest-ratio instance exists, and we can show that there are largest-ratio instances with many critical triangles.

LEMMA A.1. *There exists a largest-ratio instance, where every vertex participates in at least two critical triangles.*

PROOF. Let P be a largest-ratio instance with the minimum number of vertices, and among those with the minimum number of vertices with the maximum number of critical triangles. If a vertex A participates in no critical triangle, then we can move A while increasing area$P$ without changing the area of the largest triangle contained in P. If A participates in one critical triangle $\triangle ABC$, then there is a closed half-plane where we can move A without decreasing area$P$, and there is a line on which we can move A without changing area$\triangle ABC$. Hence there is a possibility of moving A without decreasing ratio$P$ and keeping $\triangle ABC$ critical. At some point either a vertex of P degenerates, or a new critical triangle is born; both situations contradict our choice of P. □

It follows that we may assume that a largest-ratio instance P has at least $2n/3$ critical triangles, where n is the number of vertices of P. The next three lemmas show that certain configurations of critical triangles imply a bound on ratio P.

LEMMA A.2. *Let P be a convex polygon, and let $t = \triangle ABC$ be a critical triangle with A, B, C three consecutive vertices of P in that clockwise order. Then* ratio$P \leq 2$.

PROOF. Let g be the line through A and B, and let g' be the parallel line through C. Then P has to lie in the strip between g and g'; otherwise there is a triangle larger than t in P. Similarly, for the line h through B and C, and its parallel line h' through A. Hence, P is contained in the quadrilateral Q with vertices A, B, C, and $g' \cap h'$. Since area$Q = 2$ area$t$, the lemma follows. □

Along similar lines, the next lemma can be proved.

LEMMA A.3. *Let P be a convex polygon, and let $t = \triangle ABC$ and $\triangle ACD$ be critical triangles such that the vertices A, B, C, D of P lie on the boundary of P in that clockwise order (not necessarily consecutively). Then P is a parallelogram and* ratio$P = 2$.

In other words, the lemma states that if two critical triangles share a common edge, and the respective third vertices lie on opposite sides, then the considered

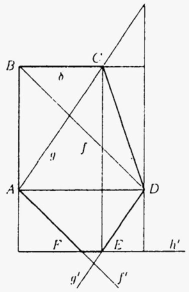

Fig. 13. The worst-case example of Lemma A.4, Case 2, with $\delta = \frac{2}{3}$.

polygon is a parallelogram. The next lemma considers the case, where the respective third vertices lie on the same side, and all involved vertices are consecutive on the boundary of the polygon.

LEMMA A.4. *Let P be a convex polygon, and let $t = \triangle ABD$ and $\triangle ACD$ be critical triangles of P with A, B, C, D four consecutive vertices of P in that clockwise order. Then* ratio$P \le 9/4$. *The ratio 9/4 is obtained by a hexagon, which is unique up to affine transformations.*

PROOF. BC has to be parallel to AD. So we may assume (by an affine transformation) that $A = (0, 0)$, $B = (0, 1)$, $D = (1, 0)$, and $C = (\delta, 1)$ with δ a parameter that varies between 0 and 1, see Figure 13. (If $\delta > 1$, then area$\triangle ABC >$ area$\triangle ABD$.) Let f be the line through B and D, and let f' be the parallel line through A. For t to be a critical triangle, P must lie completely above f'. Similarly, if g is the line through A and C, and g' is the parallel line through D, then P must lie above g'. We also want to ensure that BC is not the base of a triangle in P with area larger than $\frac{1}{2}$. This can be guaranteed iff P lies above the horizontal line h': $y = 1 - 1/\delta$. Note that the y-coordinate of $f' \cap g'$ equals $-1/(1 + \delta)$. So if $1 - 1/\delta \le -1/(1 + \delta)$, then the restriction of h' is redundant. This happens iff $\delta^2 + \delta - 1 \le 0$ which is equivalent to $\delta \le \frac{1}{2}(\sqrt{5} - 1)$ in the range $0 \le \delta \le 1$.

Case 1: $0 \le \delta \le \frac{1}{2}(\sqrt{5} - 1)$. P must be contained in the pentagon $Q = ABCDE$, where $E = f \cap g'$. The area of Q equals $\frac{1}{2}$ (for the area of $\triangle ABD$) plus $\frac{1}{2}$ (for the area of $\triangle ACD$) plus $\frac{1}{2}\delta(1 - 1/(1 + \delta))$ (for the area of $\triangle BC(f \cap g)$). Note that $\triangle ABD$ and $\triangle ACD$ "overcount" by the area of $\triangle AD(f \cap g)$ which, however, equals the area of $\triangle ADE$. Hence, area$Q = 1 + \frac{1}{2}(\delta^2/(1 + \delta)))$ which increases as δ increases for $\delta \ge 0$. So the maximum is obtained for $\delta = \frac{1}{2}(\sqrt{5} - 1)$ when area$Q = \frac{1}{2}\sqrt{5}$. Consequently, ratio$P$ = area$P$/area$t \le$ area$Q/\frac{1}{2} \le \sqrt{5} < 9/4$ which settles this case.

Case 2: $\frac{1}{2}(\sqrt{5} - 1) < \delta \le 1$. P must be contained in the hexagon $Q = ABCDEF$ with $E = g' \cap h'$ and $F = f' \cap h'$. Since the slope of f' is -1, we have $F = (1/\delta - 1, 1 - 1/\delta)$. Furthermore, C and E have the same x-coordinate δ, since

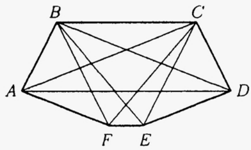

Fig. 14. A more symmetric affine transformation of the example of Figure 13.

the triangles $\triangle E(1, 1-1/\delta)D$ and $\triangle C(1, 1)(1, 1/\delta)$ are congruent. Now we can calculate the area of Q: We consider the rectangle R between the vertical lines $x = 0$ and $x = 1$ and between the horizontal lines $y = 1$ and h'. Then the area of Q equals $1/\delta$ (the area of R) minus $\frac{1}{2}(1/\delta - 1)^2$ (for the area of the triangle $\triangle AF(0, 1 - 1/\delta)$ in R not covered by Q) minus $\frac{1}{2}(1/\delta)(1-\delta)$ (for the triangles $\triangle CD(1, 1)$ and $\triangle DE(1, 1-1/\delta)$ in R not covered by Q). We obtain that $\text{area}Q = \frac{1}{2}(3/\delta - 1/\delta^2)$. The derivative $(d\,\text{area}Q)/d\delta$ equals 0 for $\delta = \frac{2}{3}$ which lies in the range of δ considered in this case. For $\delta = \frac{2}{3}$, we have $\text{area}Q = 9/8$ which is a maximum. We have shown that $\text{area}P/\text{area}t \leq 2\ \text{area}Q \leq 9/4$, which proves the inequality of the lemma in Case 2.

The above argument not only proves the upper bound of 9/4, it also demonstrates that the only possibility to attain that ratio is that P equals the hexagon Q in Case 2 for $\delta = \frac{2}{3}$ (or an affine image of this hexagon). Figure 13 shows this situation, and Figure 14 is an affine image of that same worst-case hexagon that exhibits its symmetry more clearly.

It remains to show that there is no triangle of area larger than $\frac{1}{2}$ contained in Q. Consider the cyclic sequence of triangles:

$$\triangle ABD \xrightarrow{AD \parallel BC} \triangle ACD \xrightarrow{AC \parallel DE} \triangle ACE \xrightarrow{CE \parallel AB} \triangle BCE \xrightarrow{BC \parallel EF}$$

$$\triangle BCF \xrightarrow{BF \parallel CD} \triangle BDF \xrightarrow{BD \parallel AF} \triangle BDA.$$

Each triangle is obtained from its predecessor by moving one vertex on a line parallel to the line through the two other vertices, as indicated above the arrows: We have observed before that C and E have the same x-coordinate, and so CE is parallel to AB. The fact that BF is parallel to CD follows by symmetry. The other four cases are immediate from the construction of Q. Thus, all triangles involved in this cyclic sequence have the same area$\frac{1}{2}$, and it can be easily verified that all other triangles with vertices from Q have smaller area. □

Remark. If P in the statement of the preceding lemma is a pentagon, then $\text{ratio}P \leq \sqrt{5}$ can be shown; equality holds iff P is an affine image of a regular pentagon.

Before we prove the theorem, we need the following *interspersing lemma* on which algorithms for computing largest area contained triangles are based (see [4]).

Lemma A.5. *Let P be a convex polygon, and let t and t' be critical triangles. The vertices of t and t' which are not common to both triangles alternate on the boundary of P.*

Thus, if two critical triangles of a hexagon do not share a vertex, then they must be the two alternating triangles.

Theorem A.6. $\text{ratio}P \leq 9/4$ *for every convex polygon P with at most six vertices; the bound is tight, and it is achieved by a hexagon which is unique up to affine transformation.*

Proof. Let P be a largest-ratio instance, where every vertex participates in at least two critical triangles. We know that $\text{ratio}P \geq 9/4$; so P must have at least five vertices. From Lemma A.2 it follows that at most one edge of a critical triangle is also an edge of P; so at least two edges of a critical triangle must be chords of P, i.e., line segments connecting nonadjacent vertices of P.

Case 1: P has five vertices. Then P must have at least $2 \cdot 5/3$, i.e., at least four critical triangles. Since each critical triangle has two chords of P as its edges, there must be a chord AB that participates in two critical triangles with vertices C and C'. C and C' must lie on the same side of AB (see Lemma A.3), and the assumptions of Lemma A.4 are satisfied, which shows $\text{ratio}P < 9/4$.

So P has to be a hexagon $AB'CA'BC'$, vertices in clockwise order. There are at least $(2 \cdot 6)/3 = 4$ critical triangles. If one of the diagonals AA', BB', or CC' participates in two critical triangles, then Lemmas A.3 and A.4 immediately prove the theorem. So we assume that a diagonal participates in at most one critical triangle. Now there are at most five critical triangles, two alternating and three diagonal; we conclude that there is at least one alternating critical triangle.

Case 2: P is a hexagon, and it has exactly one alternating critical triangle. Without loss of generality let this alternating critical triangle be $\triangle ABC$. Note now that every vertex participates in exactly two critical triangles, since there are only four critical triangles. So A participates in the alternating critical triangle $\triangle ABC$, and in the diagonal critical triangle that uses the diagonal AA'; similarly, for B and C. It follows that the diagonal critical triangle using AA' is either $\triangle AA'B'$ or $\triangle AA'C'$; say, without loss of generality, it is $\triangle AA'B'$. Then the two other diagonal triangles have to be $\triangle BB'C'$ and $\triangle CC'A'$, but now $\triangle A'CC'$ and $\triangle A'B'A$ contradict Lemma A.5, which settles Case 2.

The final case leads us once more into analytic calculations.

Case 3: P is a hexagon with two alternating critical triangles. If P has only two diagonal critical triangles, they cannot share a common vertex, and we get a contradiction to Lemma A.5. Hence, there are three diagonal critical triangles, and so two must share edges with the same alternating critical triangle; say with $\triangle ABC$, and the shared edges are AC and BC (the two edges have to be different, since no edge forms critical triangles with three vertices on the same side). Now these diagonal critical triangles are fixed to be $\triangle ACA'$ and $\triangle BCB'$; other possibilities are excluded by Lemma A.3 or A.5. The third diagonal critical triangle cannot be $\triangle CC'B$ or $\triangle CC'A$, since then BC (or AC) would be an edge of three critical triangles. The two remaining vertices A' and B' are symmetric, so without loss of generality let us assume that $\triangle CC'B'$ is critical. Note now that because of critical triangles $\triangle B'CB$ and $\triangle B'CC'$, $B'C$ and $C'B$ have to be parallel, and because of $\triangle B'C'A'$ and $\triangle B'C'C$, $B'C'$ and $A'C$ have to be parallel. So the position of C' is already determined by the five remaining vertices. We are free to choose A, B, and C in fixed positions, so we have to vary only A' and B' and investigate the position which maximizes ratio$P$.

We let $A = (-1, 0)$, $B = (1, 0)$, $C = (0, 1)$, $A' = (\gamma + 1, \gamma)$, and $B' = (-(\delta + 1), \delta)$ for some $0 < \gamma \leq 1$ and $0 < \delta \leq 1$. Now

$$C' = \begin{pmatrix} -(\delta + 1) \\ \delta \end{pmatrix} + \mu \begin{pmatrix} \gamma + 1 \\ \gamma - 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \eta \begin{pmatrix} -(\delta + 1) \\ \delta - 1 \end{pmatrix}$$

for appropriate μ and η. Any instance generated in this way has the property that our five critical triangles have the same area. We eliminate μ and get $\eta = 1 + \gamma/(\delta\gamma - 1)$. Therefore, the y-coordinate of C' equals

$$\eta(\delta - 1) = -\left(-\delta + \frac{1 - \gamma}{1 - \delta\gamma}\right).$$

The area of P now equals 1 (for $\triangle ABC$) plus $\frac{1}{2}\sqrt{2}(\delta\sqrt{2})$ (for $\Delta ACB'$) plus $\frac{1}{2}\sqrt{2}(\gamma\sqrt{2})$ (for $\triangle CBA'$) plus $\frac{1}{2}2(-\delta + (1 - \gamma)/(1 - \delta\gamma)$ (for $\triangle ABC'$) which gives $1 + \gamma + (1 - \gamma)/(1 - \delta\gamma)$. As is easily seen, this area is increasing as δ increases. Since P is a largest-ratio instance, there must be another critical triangle, which contradicts our assumption that there are only five critical triangles.

Consequently, the only largest-ratio instances are those constructed in the proof of Lemma A.4. □

References

[1] A. Aggarwal, J. S. Chang, and C. K. Yap, Minimum area circumscribing polygons, *Visual Comput. Internat. J. Comput. Graphics* **1** (1985), 112–117.

[2] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra, and C. Uhrig, Approximate motion planning and the complexity of the boundary of the union of simple geometric figures, in *Proc. Sixth Annual Symposium on Computational Geometry*, 1990, pp. 281–289.

[3] B. Chazelle, The polygon containment problem, *Adv. Comput. Res.* **1** (1983), 1–33.

[4] D. Dobkin and L. Snyder, On a general method for maximizing and minimizing among certain geometric problems, in *Proc. 20th Annual IEEE Symposium on Foundations of Computer Science*, 1979, pp. 9–17.

[5] L. Fejes Tóth, *Lagerungen in der Ebene, auf der Kugel und im Raum*, Springer-Verlag, Berlin, 1953.

[6] R. Fleischer, K. Mehlhorn, G. Rote, E. Welzl, and C. Yap, On simultaneous inner and outer approximation of shapes, in *Proc. Sixth Annual Symposium on Computational Geometry*, 1990, pp. 216–224.

[7] R. Frank, Private communication (1990).
[8] P. M. Gruber, Approximation of convex bodies, in *Convexity and Its Applications*, eds. P. M. Gruber and J. M. Wills, Birkhäuser-Verlag, Basel, 1983, pp. 131–162.
[9] F. John, Extremum problems with inequalities as subsidiary conditions, in *Studies and Essays Presented to R. Courant on His 60th Birthday*, Interscience, New York, 1948, pp. 187–204.
[10] R. Kannan, L. Lovász, and H. E. Scarf, The shapes of polyhedra, *Math. Oper. Res.* **15** (1990), 364–380.
[11] K. Leichtweiß, Über die affine Exzentrizität konvexer Körper, *Arch. Math.* **10** (1959), 187–199.
[12] A. Saam, Private communication (1989).
[13] O. Schwarzkopf, U. Fuchs, G. Rote, and E. Welzl, Approximation of convex figures by pairs of rectangles, in *Proc. Seventh Annual Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Vol. 415, Springer-Verlag, Berlin, 1990, pp. 240–249.
[14] C. K. Yap, Algorithmic motion planning, in *Advances in Robotics*, Vol. 1, eds. J. T. Schwartz and C. K. Yap, Erlbaum, Hillsdale, NJ, 1987, Chapter 3.