Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.

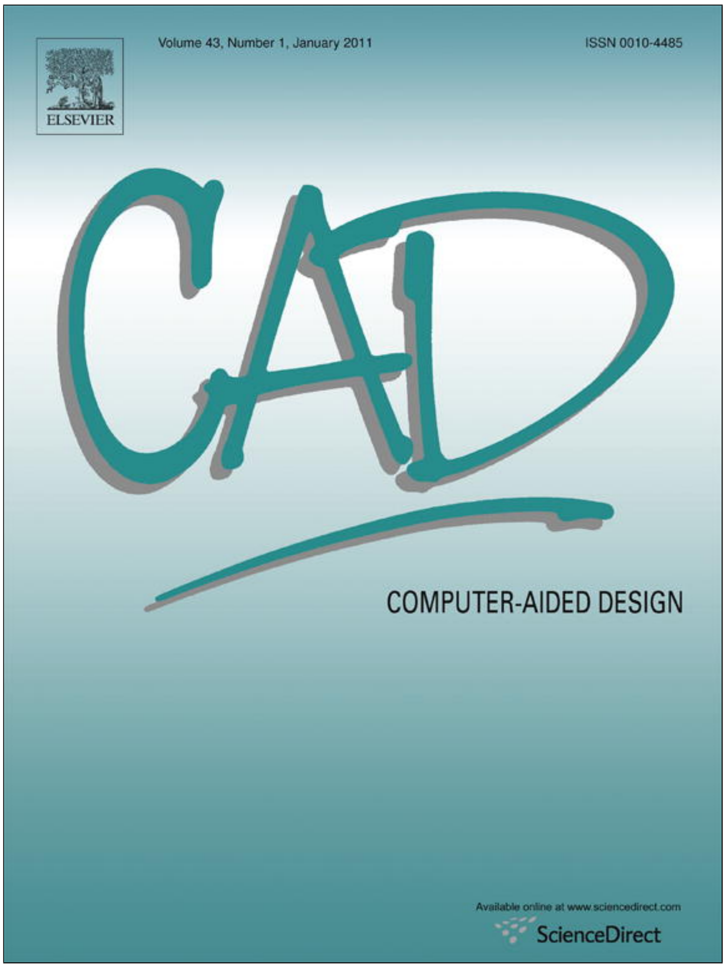

(This is a sample cover image for this issue. The actual cover is not yet available at this time.)

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Author's personal copy

Computer-Aided Design 43 (2011) 948–956

Contents lists available at ScienceDirect

Computer-Aided Design

journal homepage: www.elsevier.com/locate/cad

Corner cutting with trapezoidal augmentation for area-preserving smoothing of polygons and polylines

Dan Gordon

Department of Computer Science, University of Haifa, Haifa 31905, Israel

ARTICLE INFO

Article history:
Received 12 December 2010
Accepted 14 April 2011

Keywords:
Area preservation
Augmentation
Bounding hull
CAD
CCA
CAGD
Convexity preservation
Corner cutting
Polygon smoothing
Polyline smoothing
Subdivision

ABSTRACT

Recently, a new subdivision method was introduced by the author for smoothing polygons and polylines while preserving the enclosed area [Gordon D. Corner cutting and augmentation: an area-preserving method for smoothing polygons and polylines. Computer Aided Geometric Design 2010; 27(7):551–62]. The new technique, called "corner cutting and augmentation" (CCA), operates by cutting corners with line segments and adding the cut area of each corner to two augmenting structures constructed on the two incident edges; this operation can be iterated as needed. Area is preserved in a local sense, meaning that when a corner is cut, the cut area is added to the other side of the line in immediate proximity to the cut corner. Thus, CCA is also applicable to self-intersecting polygons and polylines, and it enables local control. CCA was originally developed with triangular augmentation, which was called CCA1. This work presents CCA2, in which the augmenting structures are trapezoids. A theoretical result from previous work is used to show that certain implementation restrictions guarantee the existence and the G^1-continuity of the limit curve of CCA2, and also the preservation of convexity. The main difference between CCA1 and CCA2 is that the limit curve of CCA1 does not contain straight line segments, while CCA2 can contain such segments. CCA2 allows the user to determine how closely each iteration follows its previous polygon. Potential applications include computer aided geometric design, an alternative to spline approximation, an aid to artistic design, and a possible alternative to multiresolution curves.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

A major topic in computer aided design (CAD) is the design of a smooth curve that interpolates or approximates a given sequence of points. The sequence may be open, i.e., it forms a polyline, or it may be closed cyclically, forming a polygon. This topic has been researched very widely—see for example Farin et al. [1].

Recent years have seen a growing interest in combining the creation of smooth curves together with the preservation of certain geometric properties, such as area preservation. Sapiro and Tannenbaum [2] and Olver, Sapiro and Tannenbaum [3], study the problem of preserving volumes, areas and lengths under smoothing operations. They show that the geometric heat flow equation can be used to obtain smoothing without shrinkage. Elber [4] combines multiresolution control and linear constraint satisfaction within the framework of multiresolution curve editing. He shows that preservation of area can be represented as a linear constraint, so this property can also be preserved. Hahmann et al. [5], by using a wavelet decomposition, also achieve area preservation, together with some other properties, such as level-of-detail (LOD) display and progressive transmission. More recently, Sauvage et al. [6] extended these two works to 3D B-spline surfaces with volume preservation by expressing the volume in trilinear form, enabling efficient LOD editing.

One of the fundamental techniques used for smoothing curves is that of *corner cutting*; see de Boor [7]. In the context of preserving various geometric properties, Mainar and Peña [8] introduce a corner cutting algorithm, called a B-algorithm, which satisfies several important properties, such as the subdivision property. However, they are not concerned with the issue of area preservation.

Smoothing of polygonal curves has also been studied in the context of computer graphics. When a smooth object is approximated by a polyhedral object, the silhouette always appears polygonal, even though various shading techniques can smooth the interior. Wang et al. [9] present a method which smooths the polygonal silhouette by approximating it, in 2D, with cubic Hermite splines. The Hermite curves are projected back to 3D and the relevant polyhedral objects are remeshed. When the remeshed objects are displayed, their contours appear smooth. However, as is evident from [9, Fig. 4], the projected area of the smoothed object can be much larger than the projected area of the original object.

Recently, Gordon [10] has presented a new approach, called "corner cutting and augmentation" (CCA), to the smoothing and area preservation problems. Starting with a low resolution polygon or polyline, CCA progressively refines it, while maintaining

E-mail address: gordon@cs.haifa.ac.il.

0010-4485/$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cad.2011.04.007

Author's personal copy

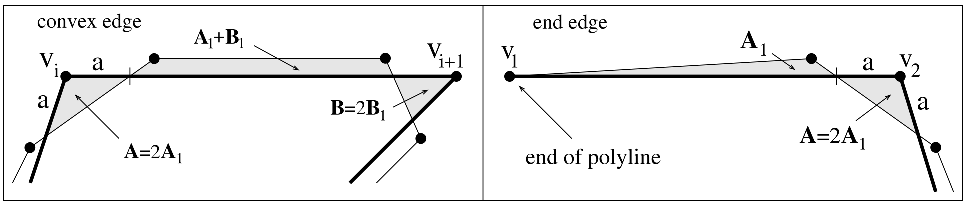

Fig. 1. Left (convex edge): the cut areas **A** and **B** are divided into two equal subareas which are added to augmenting trapezoids on the incident edges. Right (end edge): half the cut area is added to an augmenting triangle on the end edge.

a constant area. CCA uses corner cutting together with an augmentation step in which the cut areas are added back to the polygon by constructing certain augmenting structures on the edges. Area is preserved in a local sense, meaning that whenever a corner is cut, the cut area is added to the other side of the line in immediate proximity to the cut corner. This approach can be applied to simple and to self-intersecting polygons and polylines. In [10], the augmenting structures were triangles, and it was shown that certain implementation restrictions guarantee that the limit curve of CCA1 is G^1-continuous and preserves convexity.

This paper introduces CCA2, in which the augmentation structures are trapezoids. Using a theoretical result from [10], it is shown that under certain implementation restrictions, CCA2 is also G^1-continuous and preserves convexity. The main difference between CCA1 and CCA2 is that the limit curve of CCA1 does not contain straight line segments, while CCA2 can contain such segments. Technically, CCA2 is a dual binary, nonstationary subdivision scheme.

Related issues in CAGD are those of multiresolution and subdivision. [4–6] preserve area in the context of multiresolution, meaning that they maintain the same area when going from high resolution to low resolution, whereas CCA goes in the opposite direction. Subdivision schemes have a long history in CAGD; see for example the extensive review of [11] or the more recent review of [12].

CCA has the following properties and potential applications:

- It is very simple and allows local control.
- Certain simple implementation restrictions ensure the existence and G^1-continuity of the limit curve, and the preservation of convexity.
- CCA can be used for mechanical design applications which require area preservation, e.g., the design of cross-sections of complicated tubes with a specified area.
- CCA can be used in artistic design, where area preservation is a positive aesthetic feature.
- For some applications, CCA can be a simple alternative to multiresolution curves for purposes of LOD rendering and transmission, since only the initial polygon needs to be transmitted, and it can be smoothed to any level by the receiving agent.
- CCA has the potential of replacing the above-mentioned technique of Wang et al. [9] so that the smoothed object will have the same projected area as the original objects.

The two versions of CCA are suitable for different applications. A typical situation in which CCA2 is preferable to CCA1 is the design of an object which is required to rest on a flat surface. This means that a cross-section of such an object should have a straight line segment at the bottom. This can be achieved automatically with CCA2, because in the limit, every edge which is longer than its two neighboring edges gives rise to a straight line segment of the limit curve. Furthermore, in such an application, the bottom part of the cross-section should be locally convex, and in this case, the straight line segment of the limit curve is parallel to the original long edge. This property contributes to the simplicity of designing with CCA2. Another advantage of CCA2 is that certain implementation restrictions ensure that when it is applied to a regular polygon, the result will also be regular; this is not possible with CCA1.

The rest of the paper is organized as follows. Section 2 provides a general description of CCA2, and presents a result from [10] which is used in Section 3 to prove the convergence properties of CCA2. Section 4 and Section 5 develop certain design issues, while Section 6 lists some properties and examples of CCA2. Section 7 summarizes the paper and suggests some topics for further research.

2. Description of CCA2 and background

2.1. Outline of CCA2

Assume that our polygon or polyline is given by a sequence of vertices $v_1, \ldots, v_n$, with an edge joining every pair (v_i, v_{i+1}); in a polygon, there is also an edge joining (v_n, v_1), and all index operations are cyclic in n. Let v_i be a corner vertex (i.e., not the start or end of a polyline). For some $0 < \alpha < 0.5$, we mark a point on each adjacent edge whose distance from v_i is $a = \alpha \times$ [length of the shorter edge]; the corner at v_i is then "cut" by joining these points.

In the figures below, a denotes the length of the cut part of an edge, and areas are denoted by boldface letters. Note also that some figures may be distorted due to illustrative requirements. The area of a cut corner at v_i is compensated for in equal parts by the augmenting structures on the edges incident with v_i. Furthermore, the cut corners are isosceles triangles. Fig. 1 shows the construction for convex and end edges. In a convex edge, the two areas cut from the corners, denoted by **A** and **B**, are divided into two equal subareas, $\mathbf{A} = 2\mathbf{A}_1$ and $\mathbf{B} = 2\mathbf{B}_1$. To compensate for the cut areas, an augmenting trapezoid, of area $\mathbf{A}_1 + \mathbf{B}_1$, is constructed on the edge; and the other two subareas are allotted to the two other edges incident with v_i and v_{i+1}. The trapezoid is constructed on the opposite side to that of the two cut areas. The sides of the trapezoid are continuations of the line segments which cut the corners.

In the case of a polyline, there are two possibilities for handling the end edges, depending on design specifications. One possibility is shown in the right part of Fig. 1. A triangle is constructed on one side so as to compensate for half the area of the cut corner. The base of the triangle is the entire end edge (without the cut part), and one side of the triangle is a continuation of the corner cutting line.

In the limit, the above approach will result in a tangent at the end point which has a different slope from the end edge. If required, this approach can be modified so that the slope is maintained. This is done by cutting the corner with a triangle that is half the size of a regular-sized triangle, without constructing an augmenting triangle. The entire area of the cut triangle will then be used for the augmenting trapezoid on the neighboring edge.

If the two cut areas are on different sides of the line, then the edge is called an "inflection edge". Such an edge is handled at the first iteration differently from the successive iterations. The reason for this is the need to maintain local control and to guarantee the G^1-continuity of the limit curve, as will be detailed later. In the first

Author's personal copy

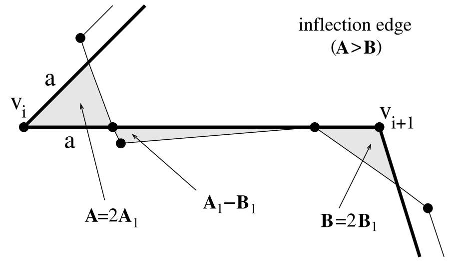

Fig. 2. Inflection edge at the first iteration, assuming area **A** is larger than area **B**: the augmenting triangle compensates for half the difference in the cut areas.

iteration, the inflection edge is handled by CCA2 as shown in Fig. 2: denote by $\mathbf{A}_1$ and $\mathbf{B}_1$ the two halves of the cut areas. The area of the augmenting triangle is $|\mathbf{A}_1 - \mathbf{B}_1|$, and it is constructed on one of the sides of the inflection edge so as to compensate for the difference in the cut areas. The figure shows the case when $\mathbf{A} > \mathbf{B}$. If $\mathbf{A} < \mathbf{B}$, then the augmenting triangle is constructed on the other side of the inflection edge. Note that one side of the triangle is a continuation of the corner cutting line. If $\mathbf{A}_1 = \mathbf{B}_1$, no triangle is constructed. This construction allows us to maintain local control.

This mode can also be used in successive iterations, but our proof of G^1-continuity of non-convex limit curves requires that inflection edges should, in the limit, maintain a fixed slope. Therefore, starting from the second iteration, inflection edges are handled as follows: if, in Fig. 2, $\mathbf{A}_1 \neq \mathbf{B}_1$, then the larger cut corner is decreased so that the two cut corners become equal (and both are isosceles). This way, no augmenting structure is constructed. The reason this approach is not used in the first iteration is that there may be several successive inflection edges, which would mean that all their cut corners should have equal areas, and this would spoil the property of local control. It is easy to see that after the first iteration, inflection edges become isolated, i.e., there are no two adjacent inflection edges. This enables local control—see Section 6.

Our implementation raises the question of how α should be chosen. Some restrictions on α are necessary in order to ensure G^1-continuity and the preservation of convexity. Other than that, the choice of α is a design issue—it can be constant, or it can depend on the angle at the vertex. α determines the extent to which the smooth curve follows the original polygon. A relation between α and the angle will be called an "α-profile", and Section 4 presents a particular profile which is based on the idea that regular polygons should produce regular polygons.

2.2. Convex polygons and their bounding hulls

In this subsection, we describe how, given a convex polygon P, we can construct a certain bounding convex polygon Q, called the bounding hull of P. This concept, apparently considered for the first time in [10], was used to prove the convergence properties of CCA1, and it will also be used for CCA2. We denote the boundary of a set S in the plane by ∂S, and its interior by $\text{int}\, S = S - \partial S$. We assume that all polygons are finite and have a non-empty interior.

Definition 1. Let P be a convex polygon with all internal angles $>\pi/2$. The *bounding hull*, $Q = \mathbb{B}(P)$, is defined as follows: consider an edge AB of P, as in Fig. 3, and let C be the intersection of the extensions of the two edges neighboring AB. Q is formed by connecting all such vertices in the same order as their corresponding edges in P.

$\mathbb{B}(\cdot)$ will also be used to represent the relation between vertices and edges of P and their corresponding edges and vertices in $\mathbb{B}(P)$, e.g., in Fig. 3, $C = \mathbb{B}(AB)$ and $CD = \mathbb{B}(B)$. Note that in Definition 1, C and P are on opposite sides of the line defined by AB because the internal angles at A and B are $>\pi/2$. The following result was proved in [10, Theorem 1].

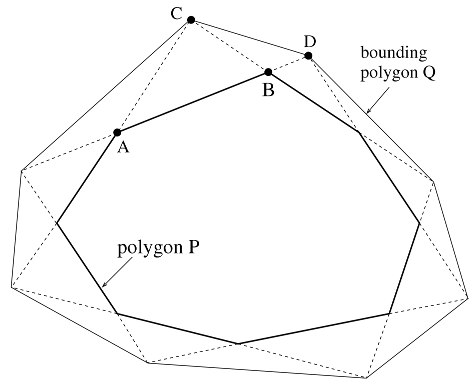

Fig. 3. The bounding hull: $Q = \mathbb{B}(P)$.

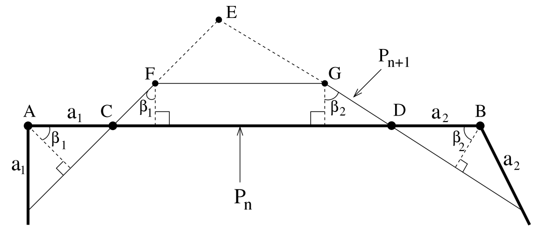

Fig. 4. Construction of the augmenting trapezoid.

Theorem 1. *Let $\{P_n\}_{n\geq 0}$ be a sequence of convex polygons such that all angles of P_n are $>\pi/2$ and all angles of P_n go to π as $n \to \infty$. Assume also that for some $\epsilon > 0$, the area of $P_n > \epsilon$. Let $Q_n = \mathbb{B}(P_n)$, and assume that $Q_{n+1} \subset \text{int}(Q_n)$. Then the sequences $\{P_n\}$ and $\{Q_n\}$ both converge to the same continuous and convex curve P^*.*

The proof of Theorem 1 proceeds along the following lines:

- If P is convex with angles $>\pi/2$, and $Q = \mathbb{B}(P)$, then Q is convex and $P \subset \text{int}Q$.
- The Hausdorff distance between P and Q, denoted by $h(P, Q)$, is always obtained as the (regular) distance between a vertex of Q and its corresponding edge in P. For example, in Fig. 3, such a distance would be between C and its orthogonal projection on AB. It is well known that the Hausdorff distance between two sets is zero iff the sets are equal.
- If S is a bounded convex set in the plane with a non-empty interior, then ∂S is a continuous curve. This is actually a well-known result, and many texts define a convex curve in the plane as the boundary of such a set.
- The main proof defines $P^* = \partial(\bigcap_{n=1}^{\infty} Q_n)$, and shows that $\{P_n\}$ and $\{Q_n\}$ both converge to P^* by proving that $h(P_n, Q_n) \underset{n\to\infty}{\longrightarrow} 0$.

3. Smoothness and convexity preservation of CCA2

3.1. Convex polygons

In this subsection we assume that the initial polygon, P_0 is convex. In the following, n denotes the iteration number. Fig. 4 provides the details for CCA2. AB is an edge of P_n, with corners cut at A and B. β_1, β_2 denote the half-angles of the cut corners, and a_1, a_2 are the cut lengths, so $a_i \leq \alpha|AB|$. The corner cutting lines intersect AB at C and D, meeting at E. The augmenting trapezoid, $FCDG$, is formed by passing a straight line parallel to AB and cutting the corner cutting lines at F and G.

Definition 2. CCA2 is required to satisfy the following conditions:

Author's personal copy

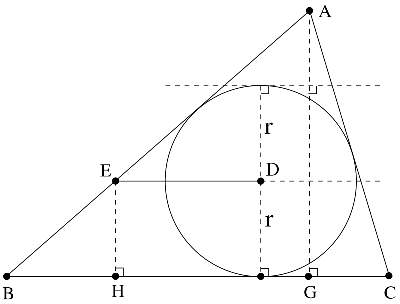

Fig. 5. Illustration of Lemma 2.

1. The area of $\triangle ECD$ is strictly greater than half the sum of the cut areas of the corners.
2. For some $\epsilon \in (0, 1)$, $\epsilon < \alpha < (1 - \epsilon)/3$, for all angles $>\pi/2$.
3. The area of trapezoid $FCDG$ is equal to half the areas cut at corners A and B.
4. Each side of the augmenting trapezoid is smaller than or equal to its adjacent cut segment, i.e., $|CF| \le |AC|$ and $|GD| \le |DB|$.

Note that each triangle at the side of the trapezoid is similar to half the cut triangle at the same side. It follows that condition 4 is equivalent to the condition that the area of each such triangle is less than or equal to half the area of the cut triangle.

Let $\{P_n\}_{n\ge1}$ be the sequence of polygons obtained by successive applications of CCA2, and denote $Q_n = \mathbb{B}(P_n)$. If the limit curve exists, we will denote it by P^*. Our main result for CCA2 is the following:

Theorem 2. *Under the above conditions, P^* exists, it is convex and G^1-continuous.*

The proof will will be a consequence of the following intermediate results:

- All internal angles tend to π as $n \to \infty$.
- The Q_ns are nested.
- From Theorem 1 and the above, it follows that the sequences $\{P_n\}$ and $\{Q_n\}$ converge to the same continuous convex curve P^*.
- Every vertex of P_n is external to P^*, and P^* intersects every edge of P_n exactly twice.
- P^* is G^1-continuous.

Lemma 1. *All angles approach π at an exponential rate as $n \to \infty$, and every P_n is convex.*

Proof. Let P_n be the polygon obtained after n iterations, as shown in Fig. 4. It is easy to see from the figure that the two angles marked with β_1 are equal (and the same holds for β_2). Hence, the internal angles formed at the top of the trapezoid are $\beta_i + \pi/2$. This means that one iteration of CCA2 replaces the angle $2\beta_i$ by $\beta_i + \pi/2$. It is clear from this that if the original polygon is convex, then this property is preserved.

To see why the angles approach π at an exponential rate, denote, in Fig. 4, $\gamma_1 = 2\beta_1$, $\gamma_2 = \beta_1 + \pi/2 = (\gamma_1 + \pi)/2$, and in general, denote by $\gamma_k = (\gamma_{k-1} + \pi)/2$ the angle obtained after k iterations. From this, it is straightforward to derive $\pi - \gamma_k = (\pi - \gamma_1)/2^{k-1}$. □

The following simple lemma will be used in the proof that the Q_ns are nested.

Lemma 2. *Let ABC be a triangle, and D the center of its incircle, as shown in Fig. 5. Consider a line through D, parallel to BC, and cutting AB at E. Then $|EB| < |AB|/2$.*

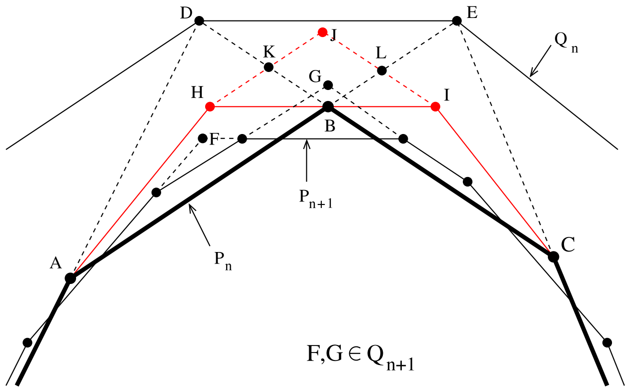

Fig. 6. Explanation of $Q_{n+1} \subset \text{int}(Q_n)$.

Proof. Let r be the radius of the incircle, and G, H be as shown. Clearly, $|EH| = r < |AG|/2$, and the result follows from the similarity of $\triangle ABG$ and $\triangle EBH$. □

We now have the following result about the sequence of bounding polygons $\{Q_n\}$:

Lemma 3. $Q_{n+1} \subset \text{int}(Q_n)$.

Proof. Fig. 6 shows a section ABC of P_n, an edge DE of Q_n, the construction of P_{n+1}, and two points, F, G belonging to Q_{n+1}. The points of Q_{n+1} are of two types: F is the intersection of the extensions of two corner cutting lines, and G is the intersection of the extensions of two lines which are the tops of two trapezoids. In Fig. 6, consider the three corner cutting lines at A, B, C shifted (parallel to their original direction) to their zero-cut limit, i.e., passing through A, B, C, respectively. The resulting lines are shown as AH, HI, IC in Fig. 6. Since the corner cutting lines form isosceles triangles, these new lines are angle bisectors: AH bisects $\angle DAB$, CI bisects $\angle ECB$, and HI bisects both $\angle DBA$ and $\angle EBC$. In $\triangle DAB$, the two angle bisectors meet at H, which is internal to $\triangle DAB$. Therefore, F is also internal to $\triangle DAB$.

To show that G is internal to $\triangle DBE$, consider the two lines through H and I parallel to the trapezoid tops meeting at G. These two lines are also parallel to AB and BC. We shall show that even the intersection of these shifted lines, denoted by J, is internal to $\triangle DBE$. Let K and L be the intersections of HJ and IJ with BD and BE, respectively. Note that $JKBL$ is a parallelogram with sides parallel to AE and CD.

Since AH, HI, IC are angle bisectors, H and I are the centers of the incircles of $\triangle DAB$ and $\triangle EBC$, respectively. By Lemma 2, $|BK| < |BD|/2$ and $|BL| < |BE|/2$, so vertex J of parallelogram $JKBL$ is strictly interior to $\triangle DBE$. □

Corollary 1. *Both sequences $\{P_n\}$ and $\{Q_n\}$ converge to the same continuous and convex limit curve P^*.*

Proof. Immediate from Theorem 1. □

Lemma 4. *Every vertex of P_n is external to P^*.*

Proof. The proof is illustrated with the aid of Fig. 7, which shows a vertex A of P_n cut by an edge BC of P_{n+1} at D and E. F is the projection of A on BD. The base angles of the cut triangle ADE, which is isosceles, are denoted by β. Consider now P_{n+2} (shown in red), which cuts the vertices B and C of P_{n+1}, and forms the trapezoid $GHIJ$ on top of BC.

Let $\gamma = \angle IJG$ be a base angle of $GHIJ$; we will show that $\gamma = \beta/2$. CK is parallel to AE because it is part of the top of the trapezoid formed on the edge of P_n containing AE. The cut triangle CJK is

Author's personal copy

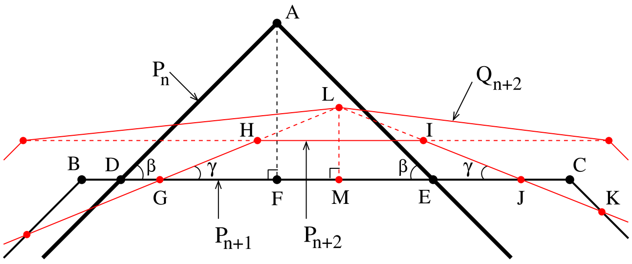

Fig. 7. Diagram for the proof that vertex A of P_n is external to P^*.

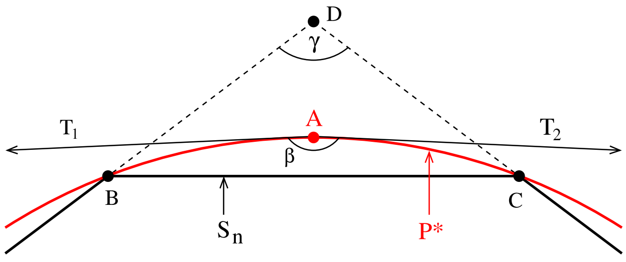

Fig. 8. Illustration of the proof that P^* is G^1-continuous.

isosceles, its base angles are equal to γ, and $\angle JCK = \pi - \beta \Rightarrow \beta = 2\gamma$. Similarly, $\angle HGJ = \beta/2 = \gamma$.

We now extend the sides GH and IJ of $GHIJ$ to meet at L, and denote by M the projection of L on BC. We will show that $|LM| < |AF|$. Assume w.l.o.g. that M lies on F or to the right of F. We will show first that $|MJ| < |AE| + |FE|$. If $M \in [F, E]$, then $|MJ| = |ME| + |EJ|$, and $|ME| \leq |FE|$. Also, by conditions 2 and 4 of Definition 2, we have $|EJ| < |EC| \leq |AE|$, and so $|MJ| < |AE| + |FE|$. If M falls to the right of E, then $|MJ| = |EJ| - |EM| < |EC| \leq |AE|$.

We now have $|MJ| < |AE| + |FE| = |AE|(1 + \cos\beta) \Rightarrow |MJ|/(1 + \cos\beta) < |AE|$. Therefore, $|LM| = |MJ| \tan\gamma = |MJ| \tan(\beta/2) = |MJ| \sin\beta/(1 + \cos\beta) < |AE| sin\beta = |AF|$.

Note that L is a vertex of Q_{n+2}, and its adjacent vertices (in Q_{n+2}) lie on the line defined by HI, as shown in Fig. 7. It follows that A is external to Q_{n+2}. Since the Q_ns are nested and converge to P^*, A is also external to P^*. □

Lemma 5. *Every edge of P_n cuts P^* at exactly two points.*

Proof. Let e be an edge of P_n. By Lemma 4, e's vertices are external to P^*. Consider now the sequence of trapezoids constructed above e by $P_{n+1}, P_{n+2}, \ldots$: their tops are parallel to e, external to P_n, and their distance from e increases monotonically. Furthermore, their distance from e is bounded, because $P_k \subset Q_k$ for all k and the Q_ks are nested. Therefore, these tops converge to a point (or a straight line segment) of P^*, which is also external to P_n. Since P^* is convex, it intersects e at exactly two points. □

Lemma 6. *The limit curve of CCA2 is G^1-continuous.*

Proof. By Lemma 5, every edge of P_n intersects P^* at exactly two points. Therefore, the intersection of P_n and P^* forms a secant polygon of P^*, which we denote by S_n. We continue the proof with the aid of Fig. 8. Assume that A is a point on P^* with two non-colinear rays of support emanating from it, T_1 and T_2, with an angle $\beta < \pi$ between them, as shown in the figure. We assume that the lines defined by T_1 and T_2 are tangents to P^*. Note that although P^* may contain straight line segments, A cannot be internal to such a segment (however, either one of P^*'s sections AB and AC can contain a straight line segment).

Since the angles of P_n go to π, it is easy to see that the same holds also for S_n. Therefore, the angles between two edges of S_n that are separated by one edge, also go to π. Assume that n is large enough so that the smallest angle between two edges of S_n separated by one edge is $> \beta$. We also assume that A does not coincide with a vertex of S_n (otherwise, we use S_{n+1}). Fig. 8 shows an edge BC of S_n and the part of P^* containing A inside the triangle DBC formed by BC and the extensions of its adjacent edges in S_n. We denote $\gamma = \angle BDC$. By our choice of n, $\gamma > \beta$.

Consider T_1: it is a supporting ray of P^*, which is convex, so it cannot strictly separate B from C because both lie on P^*. Hence, T_1 cannot cut $\triangle DBC$ at an interior point of BC. Therefore, T_1 cuts $\triangle DBC$ at BD, and similarly, T_2 cuts $\triangle DBC$ at CD. Since A is interior to $\triangle DBC$, we have $\beta > \gamma$, and this is a contradiction. □

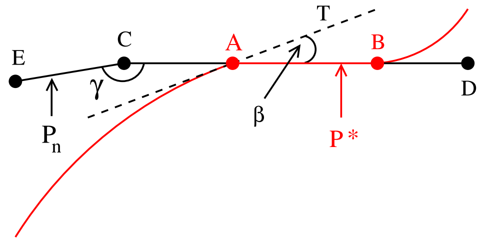

Fig. 9. Illustration of the G^1-continuity at inflection points.

Proof of Theorem 2. This follows immediately from Corollary 1 and Lemma 6. □

3.2. Non-convex polygons and polylines

We will prove that G^1-continuity also holds for non-convex polygons and polylines. Recall that starting from the second iteration, inflection edges do not change their slope and they are isolated.

Theorem 3. *The limit curve obtained by applying CCA2 to non-convex polygons and polylines is G^1-continuous at all points.*

Proof. Since we start out with a finite polygon (or polyline), there are only a finite number of inflection points in the limit curve. We call a section of a polygon or polyline convex if, by joining its endpoints, we get a convex polygon. The results of the previous section hold for convex sections, so it remains to prove G^1-continuity of the limit curve P^* at the limit points of inflection edges.

After the initial iteration, inflection edges are simply shortened at both ends by successive CCA2 iterations, and so each inflection edge gives rise to a sequence of strictly embedded edges. In the limit, there are two possible outcomes of such a sequence:

Case 1: the limit is a straight line segment of P^*, such as AB in Fig. 9.

Case 2: the limit is a single point, as if A and B coincide.

We will consider the two cases separately.

Case 1 is illustrated in Fig. 9. We will show the G^1-continuity at points such as A. Assume that there is a tangent line T passing through A and making an angle $\beta > 0$ with AB. Note that due to the isolation of inflection edges, there is a convex section of P^* attached to A, as shown in the figure. Consider now some edge CD of P_n, colinear with AB, and let EC be the adjoining edge at C. We denote $\gamma = \angle ECA$. Since the angles tend to π, we can take n to be sufficiently large so that $\gamma > \pi - \beta$. Therefore, EC cannot intersect T, and since T is a line of support of the adjoining convex section of P^*, EC cannot intersect the convex section. This contradicts Lemma 5.

Case 2 is actually quite similar to Case 1. Let A be a point on P^* which is the limit of a sequence of embedded inflection edges. Now, instead of the line segment AB of Case 1, consider a line ℓ through A containing the sequence of inflection edges converging to A. The rest of the proof is now identical to Case 1, with ℓ replacing AB. CD will now be an inflection edge belonging to the sequence converging to A and contained in ℓ. □

Author's personal copy

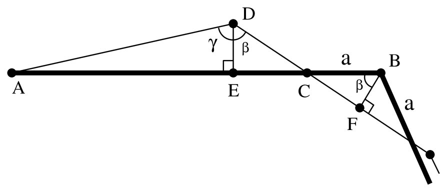

Fig. 10. The augmenting triangle at an end edge of a polyline.

3.3. End edges of polylines

We now consider end edges of a polyline, and we will show that the angle between such edges and their neighboring edge tends to π.

Lemma 7. *The angles at an end edge tend to π at an exponential rate as $n \to \infty$.*

Proof. Let β denote the half-angle between an end edge and its neighboring edge. If we use the option of not constructing an augmenting triangle on the end edge, then it is easy to see that after one iteration, the new half-angle is $(\beta + \pi/2)/2$, i.e., it is the average of β and $\pi/2$, and the result follows.

We now consider the option of constructing an augmenting triangle, and the proof proceeds with the help of Fig. 10. Let AB be an end edge, β the half-angle at B, and DAC the augmenting triangle. Clearly, $\angle CDE = \beta$. Assume w.l.o.g. that $|AB|$=1. Recall that the cut triangle at B is isosceles, and for some $\epsilon \in (0, 1)$, $\alpha < (1-\epsilon)/3$. Therefore, $|BC| < (1-\epsilon)/3$. The areas of $\triangle DAC$ and $\triangle BCF$ are equal, and so in the two similar triangles DEC and BCF, the sides of $\triangle DEC$ are smaller than their corresponding sides in $\triangle BCF$. Therefore, $|EC| < |DC| < |BC| < (1-\epsilon)/3$. Also, $|AE| = 1 - |EC| - |BC| > 1 - 2(1-\epsilon)/3$. From this we get:

$$\frac{\tan\gamma}{\tan\beta} = \frac{|AE|}{|EC|} > \left(1 - \frac{2(1-\epsilon)}{3}\right) \cdot \frac{3}{1-\epsilon} = 1 + \frac{3\epsilon}{1-\epsilon}.$$

Let $\{\beta_n\}_{n\geq 1}$ be the sequence of half-angles at the end edge, with $\beta_1 = \beta$. Also, let $\gamma_1 = \gamma$, and $\gamma_n = 2\beta_{n+1} - \beta_n$. From the above, we have $\tan\gamma_n > (1+c)\tan\beta_n$ for all n, where $c > 0$. Therefore, $\tan\gamma_n > \tan\beta_n$ for all n, and so $\beta_{n+1} = (\beta_n + \gamma_n)/2 > \beta_n$. Since the sequence $\{\beta_n\}$ is monotonically increasing and bounded, it converges to a limit, which we denote by β^*. From $\gamma_n = 2\beta_{n+1} - \beta_n$, we also have $\gamma_n \underset{n\to\infty}{\longrightarrow} \beta^*$. Claim: $\beta^* = \pi/2$. Assume otherwise, then $\tan\beta^*$ is finite, so from $\tan\gamma_n > (1+c)\tan\beta_n$ we get $\tan\beta^* \geq (1+c)\tan\beta^*$, which is impossible. The proof of the exponential rate of convergence is similar to previous such proofs. □

4. Angle-dependent choice of α

4.1. Preliminaries

As mentioned in the Introduction, the term "α-profile" denotes any dependency of α on the angle of the cut corner (a constant α is also an α-profile). The choice of an α-profile is a design issue, to be determined by the user. We describe below a certain approach to this question, which is possibly useful in a wide variety of situations, including the application to non-convex polygons and polylines. Our approach is motivated by striving to achieve, in some sense, curves that are close to arcs of a circle. This is done by considering regular polygons, because in the limit, regular polygons approach a circle.

The α-profile for CCA2 will satisfy condition 2 of Definition 2. Even if the α-profile specifies that $\alpha \to 0$ as the angle goes to zero, this does not contradict condition 2 of Definition 2, because this condition is only required for angles $>\pi/2$. Recall that we made the following assumptions about CCA2:

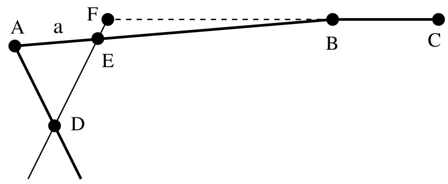

Fig. 11. An example showing that α cannot always depend only on the angle.

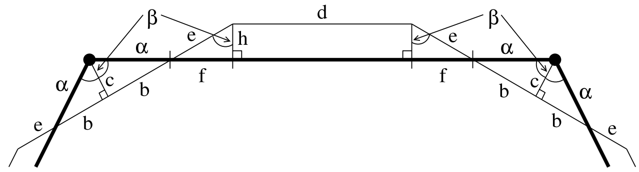

Fig. 12. Cut corners of a regular n-gon (edge length = 1).

- The cut corner is always an isosceles triangle. And the two equal sides of the triangle are taken as α times the length of the smaller edge.
- The area of a cut corner is divided equally between the two incident edges.

We will call an α-profile for CCA2 *regular* if, whenever CCA2 is applied to a regular polygon with this profile, the result is also a regular polygon. We will derive below a regular α-profile for CCA2.

It should be mentioned that in some extreme situations, it may not be possible to let α depend only on the angle. An example is shown in Fig. 11: the angle at B is very large and the extension of the edge CB meets the cutting line through D and E at a point F so that $\triangle FEB$ is too small to contain the augmenting structure. In such cases, one or both of the αs should be reduced until the augmenting structure, as required by Definition 2, can be constructed. Clearly, this can always be done by making α sufficiently small. Since the angles tend to π at an exponential rate, such cases of extreme differences between adjacent angles can only occur for a finite number of iterations. An alternative approach is detailed in Section 5.

4.2. The regular α-profile for CCA2

Fig. 12 shows an edge of a regular polygon, with all edge lengths equal to 1, on which an augmenting trapezoid is constructed. Note that in the figure, α is also the actual cut length. β denotes the half-angle of the original polygon, and it is clear that all the angles denoted by β in the figure are equal. If the original polygon contains n edges, then its internal angle is $2\beta = (n-2)\pi/n$, so $\beta = \pi/2 - \pi/n$. The new polygon contains $2n$ edges, and if it were regular, then its internal angle would be $\gamma = (2n-2)\pi/(2n) = \pi - \pi/n$. However, we can see from Fig. 12 that the internal angle of the new polygon is $\beta + \pi/2 = \pi - \pi/n$. This means that regardless of the choice of α, one iteration of CCA2 applied to a regular n-gon yields a $(2n)$-gon with the same internal angle as a regular $(2n)$-gon.

The problem that remains now is to determine α so that the resulting polygon will have edges of equal lengths. From Fig. 12, we see that this means we should have $d = 2b + 2e$. The area equality condition means that $(d+f)h = bc$. We also have $d + 2f + 2\alpha = 1$. A few simple calculations yield the following quadratic equation for α:

$$[8(1+\sin\beta)\cos^2\beta]\alpha^2 - [8(1+\sin\beta)]\alpha + [2+\sin\beta] = 0. \quad (1)$$

Author's personal copy

Table 1
Values used for creating the regular α-profile for CCA2.

Interior angle	0	40	60	90	120	150	180
No. of edges	n/a	n/a	3	4	6	12	∞
α	0	0.22311	0.25842	0.22311	0.20221	0.19103	0.18750

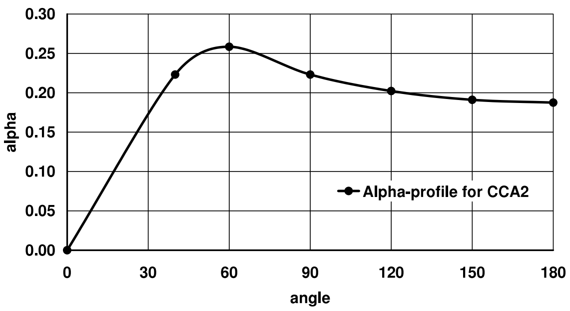

Fig. 13. The regular α-profile for CCA2.

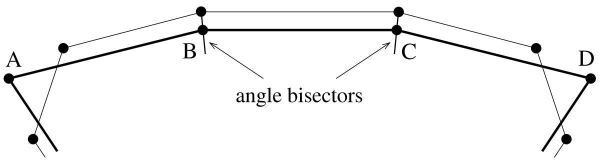

Fig. 14. The inflation technique for creating a trapezoidal chain.

It is easy to see that as $\beta \to \pi/2$, $\alpha \to 3/16 = 0.1875$. Table 1 shows the values that were used in creating the regular α-profile. As in the case of CCA1, we chose to let α go to zero as the angle goes to zero, and we added an additional control point at 40°. The line plot of the profile, shown in Fig. 13, is very similar to that of CCA1, with the maximum also at 60°.

Condition 2 of Definition 2 states that for some $\epsilon \in (0, 1)$, $\epsilon < \alpha < (1 - \epsilon)/3$ for all angles $>\pi/2$. We will see that this property is satisfied by the regular α-profile. From Table 1 and Fig. 13 we have $0.1875 < \alpha < 0.25$ for angles $>\pi/2$. So, if we take any $0 < \epsilon \leq 0.1875$, we will have $\epsilon < \alpha$. Also, we will have $(1-\epsilon)/3 \geq (1-0.1875)/3 \approx 0.27$, which means that the second inequality will also be satisfied.

5. The inflation technique for CCA

As noted in Section 4, it may not always be possible to allow α to depend only on the angle, and it may be necessary to decrease α in order to allow construction of the augmenting structure. This section describes an alternative approach for CCA2, and it can also be used for CCA1.

5.1. The inflation algorithm

The basic idea is to "inflate" a section of the polygon, as shown in Fig. 14. Consider the cut corners at A and D, and assume that the angles at B and C are too large to enable the construction of suitable trapezoids on the edges AB and CD. The inflation technique constructs a *trapezoidal chain* on the edges between A and D, as shown in the figure. The extreme sides of the trapezoids (near A and D) are continuations of the cut lines. All the other sides of the trapezoids are formed by the bisectors of the angles at B and C. The total area of the constructed trapezoids should be equal to half the sum of the areas cut from A and D. Inflation can also be used with non-convex polygons and polylines, but a trapezoidal chain can only be constructed on a convex section of the polygon or polyline.

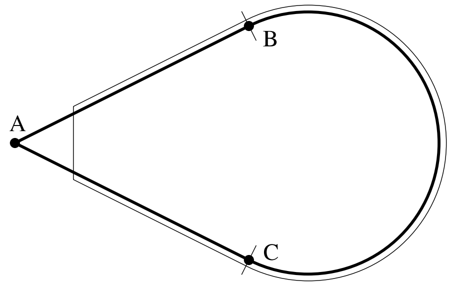

Fig. 15. Polygon inflation may be required for all edges except two (AB and AC).

The inflation method could extend to the whole polygon in extreme situations, as in Fig. 15. In such cases, one would lose the property of local control. Note that the large angles that are not cut retain their values (angles at B and C in Fig. 15), while the cut angles give rise to larger angles. From Lemma 1, an angle of ϵ increases to $(\epsilon + \pi)/2$ after one iteration of CCA2, so after k iterations it becomes $\pi - (\pi - \epsilon)/2^k$. This means that all cut angles approach π at an exponential rate, so inflation will be required only a small finite number of times.

The principle of inflation itself does not provide a method for determining which corners are cut and which corners participate in some trapezoidal chain. One method for doing that is described below. The basic idea is to start the corner cutting by going from the smallest angles to the largest, for as long as the regular trapezoid construction is possible. If all the angles have been cut, we continue to the next iteration. If at some point we can no longer construct regular trapezoids, then we construct trapezoidal chains, as detailed below.

Algorithm 1 (*The inflation technique for CCA2*)**.**

- Assume that $v_1, \ldots, v_n$ is the list of vertices in increasing angle size. If we are dealing with a polyline, then we assume that the end vertices are v_1 and v_2.
- Begin to cut corners in the order $v_1, \ldots, v_n$ (start with v_3 if dealing with a polyline).
- Whenever a cut corner v_j shares an edge with some v_i, where $i < j$, check if a trapezoid can be constructed. If yes, then construct it and continue. If v_i is an end point of a polyline, then construct the required augmenting triangle on the edge v_iv_j.
- If a trapezoid cannot be constructed on the edge v_iv_j, then cancel the cutting at vertex v_j and stop the corner cutting.
- When the corner cutting stops (for the above reason), construct trapezoidal chains between every pair of cut corners (as in Figs. 14 or 15).
- Continue to the next iteration, or exit if the iteration limit has been reached.

It is clear that if the original polygon is convex then the inflation technique will keep it convex. Furthermore, if for some segment of a non-convex polygon or polyline all the angles which are $<\pi$ are on the same side of the segment, then that property will also be preserved. Note also that the inflation method can handle the case where three sequential vertices lie on a straight line (the angle bisectors in this case are perpendicular to the common line, and the trapezoidal chain built on the colinear edges actually forms one large trapezoid).

Author's personal copy

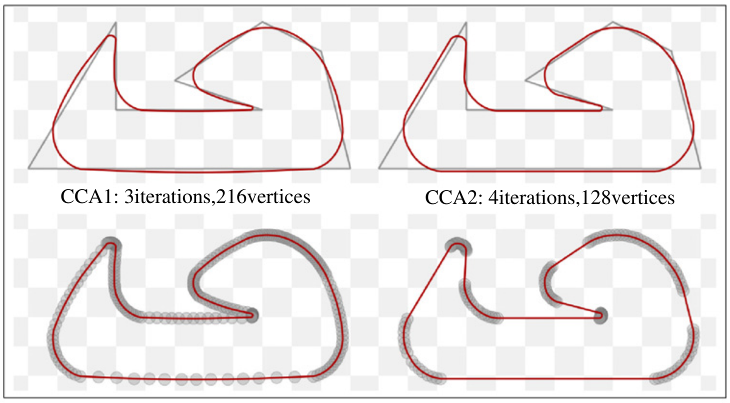

Fig. 16. A design example illustrating the differences between CCA1 and CCA2. The small circles indicate the vertices of the resulting polygons.

5.2. Implementation and analysis of Algorithm 1

Whenever a corner is cut, we get two vertices of the next iteration, and their angles are equal (and depend only on the angle of the cut vertex). Hence, for the sorted list of the next iteration, we can keep them in a batch, so the number of elements in the next list remains unchanged. Secondly, since the vertices are cut in order of increasing angle size, then the sizes of the angles of the vertices for the next iteration are also obtained in increasing angle size. Now, when the cutting process stops and the trapezoidal chains have been constructed, we are left with two sorted lists: the list of batches, and the list of vertices which have *not* been cut.

The total number of elements on the two lists is exactly n, and the sorted list for the next iteration is obtained by merging the two sorted lists. This gives us the (worst-case) time complexity of running Algorithm 1 for k iterations: it is $O(n \log n)$ for the initial sorting, and $O(kn)$ for $k - 1$ merging operations. The new vertices are always batched—two elements per batch after the first iteration, four after the second iteration, and so on. From this we can conclude that the time required by Algorithm 1 (for CCA2) for k iterations is $O(n(\log n + k))$.

It should be mentioned that inflation can also be applied to CCA1, but the implementation and analysis require too many details from [10]. We will just mention that the time required by Algorithm 1 for CCA1 for k iterations is $O(n(k + \log n)2^k)$.

6. Additional properties and results

Local control. When the position of a vertex is changed, only its adjacent edges change. This affects the neighboring angles, and so, assuming that inflation is not used, the changes propagate to the augmenting structures on the four edges closest to the vertex, but no further. This provides the user with local control. The use of inflation may disable the feature of local control.

Anchor points. Given a polygon or polyline, one can determine one or more "anchor" points as vertices which remain fixed during the smoothing operations. This is done simply by considering each sequence of vertices between two anchor points as a polyline which is smoothed independently of other such sequences. Note that this option will not, generally, result in G^1-continuity of the curve at the anchor points. The following alternative approach will rectify this: anchor points at which G^1-continuity is required are chosen as extra vertices in the interior of an edge, and the sections between anchor points are treated as polylines. However, the end edges are now handled with the option of *not* constructing an augmenting triangle on them. This way, the line slopes on both sides of the extra vertices will coincide with the edge in which the extra vertex was chosen, so the resulting curve will be G^1-continuous everywhere.

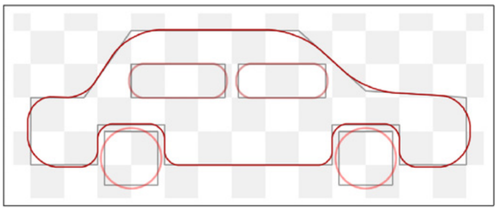

Fig. 17. CCA2 example: the original polygons (gray) and after 4 iterations (red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Sample results. Fig. 16 shows the differences between CCA1 and CCA2: with CCA1, edge lengths go to zero (as shown in [10, Lemma 5]), while CCA2 leaves some straight line segments. Edges that are longer than their neighbors will lead to straight segments in the limit curve. The top part of the figure shows the original and the smoothed polygons, while the bottom part shows the difference in the distribution of the vertices. Note that in both versions of CCA, the distribution is dense where the curvature is large and sparse where it is small. The α-profile used with CCA1 is the same one that was derived in [10] along lines that are similar to those used for the regular α-profile of CCA2.

The design goal of this example is to create a smooth (but somewhat complex) cross-section of a pipe with a given area. If the pipe needs to be laid on a flat surface, then the bottom of the cross-section should be straight, and this gives CCA2 an advantage over CCA1. The original polygon has only 8 vertices. Three iterations of CCA1 resulted in $8 \times 3^3 = 216$ vertices, and four iterations of CCA2 resulted in $8 \times 2^4 = 128$ vertices. CCA1 requires three iterations for this example because 2 iterations do not look sufficiently smooth in some sections. This is another advantage of CCA2 with this example, since significantly fewer vertices are needed to produce a smooth-looking result.

Fig. 17 shows another result with CCA2. Note that the wheels, which were designed with squares, turned into regular polygons with 128 vertices, and look like perfect circles.

An example with CCA3. An additional method, called CCA3, was also implemented. It is also based on trapezoidal augmentation and it uses the regular α-profile of CCA2, but the cut segment at the end of each edge is determined only by the value of α (in CCA2, the cut lengths are determined by the smaller edge incident with the vertex). Four iterations of CCA3 were applied to the basic polygon of Fig. 16, and the results are shown in Fig. 18. Note that this method does not approximate the original polygon as well as CCA1 or CCA2.

Author's personal copy

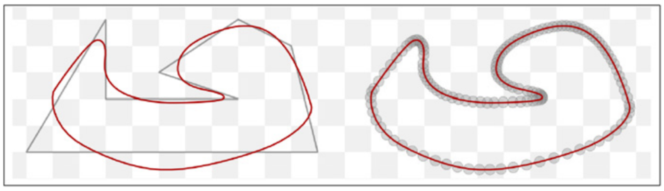

Fig. 18. Result of applying 4 iterations of CCA3 to an 8-vertex polygon, resulting in 128 vertices. Left: original and smooth polygons; right: distribution of vertices.

7. Conclusions

Corner cutting and augmentation (CCA) is a new subdivision method for smoothing polygons or polylines while maintaining the area. It is a local technique, meaning that area detracted from one side of the line is added to the other side in immediate proximity, and it allows local control. CCA operates by cutting corners and adding the cut areas by constructing augmenting structures on the edges. It is very simple and potentially useful for CAGD and artistic design. For some applications, it could replace the need for multiresolution curves for transmission purposes and LOD rendering; only the basic polygon is required, and it can be smoothed as needed.

This work presented CCA2, which uses trapezoids for augmentation, as opposed to CCA1 [10], which uses triangles. Both methods allow the user to specify an "α-profile", which determines the extent to which the smooth polygon or polyline follows the original. An α-profile for CCA2 is based on the idea of transforming a regular polygon into a regular polygon. In order to handle cases of extreme differences between the angles of the base polygon, the trapezoidal construction on one edge is extended to the construction of a trapezoidal chain on several consecutive edges.

The following are some topics for further research:

- Can the smooth curves obtained with CCA be characterized algebraically, e.g., as piecewise polynomial or rational splines?
- A related issue is the question of G^k-continuity of the limit curve for $k > 1$.
- The following inverse problem is also interesting: given a smooth curve, or a very tight polygonal approximation of such a curve, can one find a base polygon with a small number of vertices so that a few CCA iterations will produce a good approximation to the curve? A solution to this problem could provide an alternative to multiresolution curves.
- Application of Theorem 1 to other subdivision schemes. Some preliminary results have already been obtained with stationary schemes.
- Extension to 3D polyhedra: as noted in [10], the bounding hull can be extended to 3D convex polyhedra with triangular faces, but the general question of extending CCA to 3D remains open. We conjecture that Theorem 1 can be extended to higher dimensions.
- Convergence analysis of CCA3: this will be quite different from the previous methods, since these rely on the fact that cut triangles are isosceles. From various examples with CCA3, it appears that the resulting curve does not contain straight line segments, but this needs to be proved.
- Discrete curvature distributions of CCA1, CCA2 and CCA3: it would be interesting to compare these on a wide variety of test cases in order to see which method provides the smoothest curvature variation.
- Application to graphics: initial experiments indicate that CCA can be used to smooth the silhouettes of rendered polyhedral objects (and maintain the projected area). However, many details still need to be worked out.

Acknowledgments

The author would like to thank the anonymous reviewers for their careful reading of the manuscript and their many useful and detailed comments, including the question of the G^1-continuity of CCA2 applied to non-convex polygons and polylines. Thanks are due to Vlad Gorodetsky for implementing CCA2 and CCA3 on the platform which he developed with Eli Arbel for CCA1, and also to Kim Phong for implementing some changes and options to the various CCA methods.

Appendix. Supplementary data

Supplementary material related to this article can be found online at doi:10.1016/j.cad.2011.04.007.

References

[1] Farin G, Hoschek J, Kim M-S. Handbook of computer aided geometric design. Amsterdam: Elsevier; 2002.
[2] Sapiro G, Tannenbaum A. Area and length preserving geometric invariant scale-spaces. IEEE Transactions on Pattern Analysis and Machine Intelligence 1995;17(1):67–72.
[3] Olver P, Sapiro G, Tannenbaum A. Invariant geometric evolutions of surfaces and volumetric smoothing. SIAM Journal on Applied Mathematics 1997;57(1): 176–94.
[4] Elber G. Multiresolution curve editing with linear constraints. Journal of Computing and Information Science in Engineering 2001;1(4):347–55.
[5] Hahmann S, Sauvage B, Bonneau G-P. Area preserving deformation of multiresolution curves. Computer Aided Geometric Design 2005;22:349–67.
[6] Sauvage B, Hahmann S, Bonneau G-P, Elber G. Detail preserving deformation of *B*-spline surfaces with volume constraint. Computer Aided Geometric Design 2008;25:678–96.
[7] de Boor C. Cutting corners always works. Computer Aided Geometric Design 1987;4:125–31.
[8] Mainar E, Peña J. Corner cutting algorithms associated with optimal shape preserving representations. Computer Aided Geometric Design 1999;16: 883–906.
[9] Wang B, Wang W, Wu J, Sun J. Silhouette smoothing by boundary curve interpolation. In: E. Wu, H. Sun, and D. Qi (Eds.). Proc. 8th International Conf. on CAD/graphics Technology and its Applications. 2003. p. 197–202.
[10] Gordon D. Corner cutting and augmentation: an area-preserving method for smoothing polygons and polylines. Computer Aided Geometric Design 2010; 27(7):551–62.
[11] Dyn N, Levin D. Subdivision schemes in geometric modeling. Acta Numerica 2002;11:73–144.
[12] Dyn N. Linear subdivision schemes for the refinement of geometric objects. In: Proc. International Congress of Mathematicians. 2006. p. 1201–26.