Chapter 3

Perturbations of Self-Touching Configurations

3.1 Introduction

Given a self-touching configuration $\mathcal{C}$, an initial perturbed drawing D_0 which is simple, but with edges that are not necessarily straight lines, can be obtained easily:

Step 1. We start from the geometric drawing and place a circle of radius δ around every point p (δ has to be small enough such that no circle intersects another circle or an edge to which its center is not incident.)

Step 2. For each edge of multiplicity m we can easily draw m parallel line segments between the disks representing the endpoints. These segments terminate at the disk boundaries. The parallel lines must be drawn close enough to the original position to ensure that no segments coming from different edges intersect.

Step 3. Then we draw inside each disk D_p a copy of the plane graph G_p in the specification disk using the terminal points that were fixed in Step 2.

We get a plane drawing D_0 of $\mathcal{C}$ in which every vertex is at a distance at most δ from its target position. By selecting a straight-line drawing inside the disks it is possible to achieve that the edges of D_0 are polygonal chains, but this is not important for our definition. (It is not hard to see that the graphs G_p can always be drawn with straight edges.)

We are looking for a straight-line drawing D of $\mathcal{C}$ that can be obtained from this drawing D_0 by continuously deforming the edges, while keeping the graph non-crossing at all times and the vertices within the δ-disks at all times. We call such a drawing a *δ-perturbation*. In other words, a *δ-perturbation* of a self-touching configuration is a repositioning of the vertices within δ-disks consistent with the combinatorial planar embedding.

As an example, the drawing in Figure 3.1(Right) is a δ-perturbation when the dotted disks have radius δ. If 2δ is smaller than the distance between two points, this definition ensures that a vertex remains on the "correct side" of an edge which it touches, in an intuitive sense.

Note that for edges which are close to each other and almost parallel, such as in Figure 3.2, it is possible to have homotopic straight-line drawings (within δ-disks) in which points have

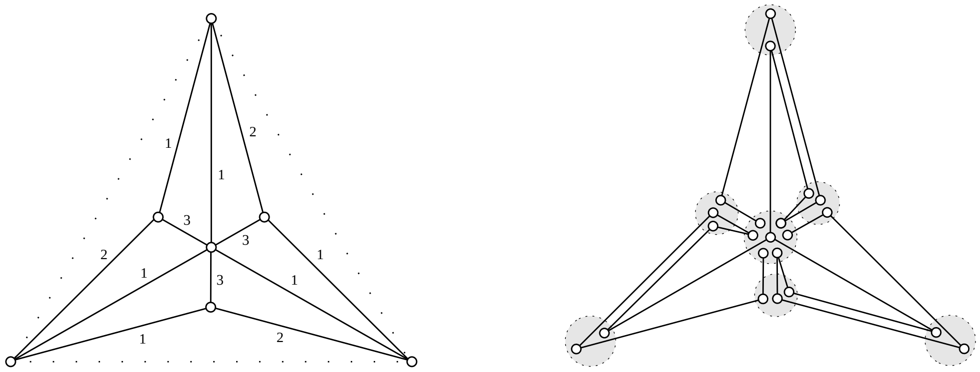

Figure 3.1: From [17]. Left: A self-touching configuration. Right: A perturbation of the self-touching configuration on the left, within the dotted circles.

switched to the other side of an edge. Hence our condition is stronger than merely requiring that D and D_0 are homotopic.

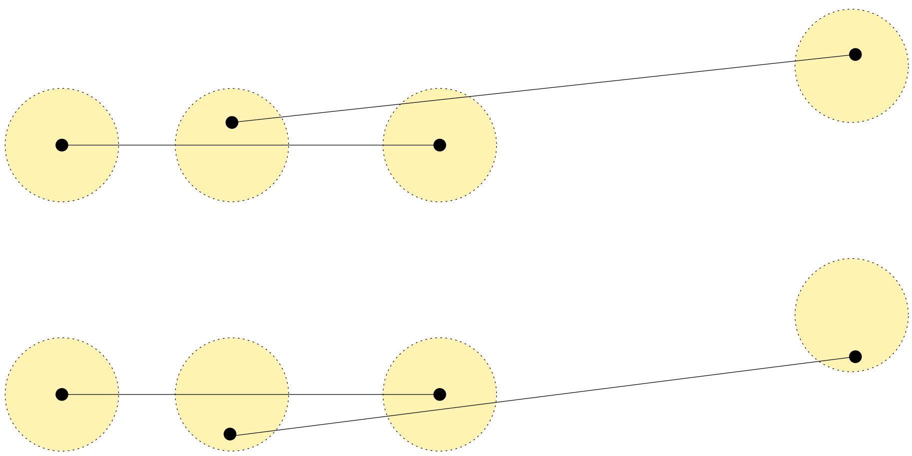

Figure 3.2: Two homotopic straight-line drawings within the shadowed disks, which are not δ-perturbations of the same self-touching configuration.

Perturbations appear in the context of locked linkages, as it has been discussed in Section 1.2. We have been able to prove the following statement, posed as a conjecture by Connelly, Demaine and Rote [17]:

Theorem 3.1. *For each self-touching configuration, and for each $\delta > 0$, there exists a δ-perturbation.*

In Section 3.2 we prove the theorem for one-dimensional self-touching configurations, that is, self-touching configurations in which all vertices lie initially on a line. We can draw any self-touching configuration as a simple planar drawing with x-monotone edges that are not necessarily straight lines, in which every vertex is at distance at most δ from its initial position, and we show that for any planar drawing with x-monotone edges, there exists a straight line embedding with given x-coordinates. In Section 3.3 we prove the theorem for the general planar case, using the one-dimensional case as a building block. In Section 3.4 we show, as a consequence of Theorem 3.1, that polygonal chains and polygonal cycles are infinitesimally flexible, which is a first step for solving the infinitesimal version of the Carpenter's Rule problem.

3.2 1-D Self Touching Configurations

In this section we prove that any one-dimensional self-touching configuration can be perturbed: we show that any one-dimensional self-touching configuration can be transformed into an x-monotone 3-connected planar graph with two x-monotone chains on the boundary, and where every intermediate vertex has left and right neighbors, and we use that any such a graph has a straight line embedding in equilibrium stress with given x-coordinates.

A curve γ is *x-monotone* if any vertical line either does not intersect γ, or it intersects γ at a single point. We say that a graph is x-monotone if all its edges are x-monotone.

Given a one-dimensional self-touching configuration, we can easily obtain a perturbed drawing which is simple and where every bar is represented by an x-monotone arc, where the arcs are not necessarily straight lines but polygonal lines, with all vertices distinct and within δ-balls. The procedure is as follows. First we draw a small disk around each vertex location and place the terminals appropriately: the terminals which have connections to the left side are placed on the left semi-circle, and similarly on the right side. Now we place the vertices in each disk: First we draw the edges that go through without touching a vertex, and the vertices that have connections both to the left side and to the right side (together with their incident edges). These edges and vertex stars have a unique order from top to bottom since they do not cross, and we draw them one after the other from top to bottom. There is no problem in drawing these components: we just place the vertex anywhere between the rightmost connection on the left and the leftmost connection on the right, below the previously drawn edges and vertices, and connect them by x-monotone paths. Finally we draw the vertices that have connections only to the left or only to the right. Such vertices may be *nested* inside each other: If a vertex v is connected to two terminals a and b on the, say, left side, and another vertex w is connected to a terminal between a and b, then w is nested within the connections of v, and we draw v before w. Since the graph is planar, the nesting relation is transitive and acyclic. Thus, there is a good order in which we can draw the vertices (together with their incident connections), beginning with the outermost levels of nesting. In the end, we just have to connect the terminals on one side of each disk to the corresponding terminals on the next disk by monotone paths, which is easy. We call such a drawing an *x-monotone representation.* Here we obtain a representation with polygonal arcs, but in the figures we draw curved arcs.

The difficulty is how to simultaneously straighten the arcs while vertices do not move outside the δ-balls. See Figure 3.3. Our strategy consists of fixing the x-coordinates in the x-monotone representation and to apply the following theorem.

Theorem 3.2. *Any planar embedding with x-monotone edges can be straightened, maintaining the same x-coordinates, with all vertices distinct.*

The straight line embedding given by this theorem is precisely the δ-perturbation. Theorem 3.2 was independently proven in 2002 by Pach and Tóth [46], but we were not aware of it.

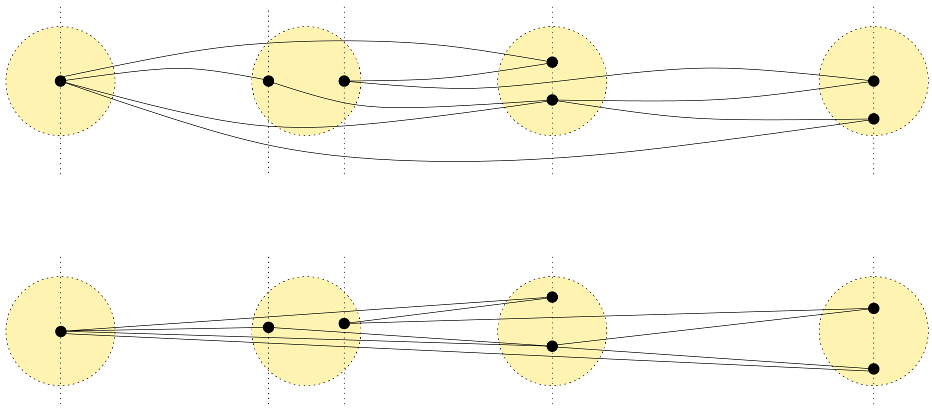

Figure 3.3: Top: An x-monotone representation of a 1D self-touching configuration. Bottom: Straight line embedding with the same x-coordinates and with vertices inside δ-balls.

For general self-touching configurations, we will need in Section 3.3 the following variation of Theorem 3.2.

Theorem 3.3. *Any planar embedding with x-monotone edges and a fixed polygonal convex outer face such that no chord overlaps any boundary bar when it is drawn straight, can be straightened maintaining the same x-coordinates, with all vertices distinct.*

To prove Theorem 3.2, we need the following lemma.

Lemma 3.1. *Any x-monotone representation of a one-dimensional self-touching configuration can be extended to an x-monotone planar graph such that:*

(a) It is 3-connected.

(b) Every interior vertex, that is, every vertex not incident to the exterior face, is adjacent to vertices on both sides of it, left and right.

(c) The boundary is an x-monotone triangle.

Proof. We perform a sequence of transformations to our x-monotone representation for obtaining the desired graph. To fulfill condition (c), we just add 3 new vertices and 3 new edges forming an x-monotone triangle on the boundary. To fulfill condition (b), for any vertex u we do the following: if u has no left neighbor, we add an x-monotone edge joining u and a vertex on its left, belonging to the same face as u and similarly for the right. Note that, after adding all these edges, each face is x-monotone, that is, every polygon has two monotone chains on the boundary. Hence we have not introduced crossings. For condition (a), we triangulate the obtained graph: add a vertex v in the interior of each non-triangular face, and connect v to each vertex of the face by x-monotone edges.

These steps are illustrated in Figure 3.4. Since within this transformation we have not introduced any loop or any multiple edge, our triangulation is 3-connected.

It is easy to achieve that the obtained graph has no vertical edges, since vertices initially connected converge to different positions and we can achieve that the new added x-monotone

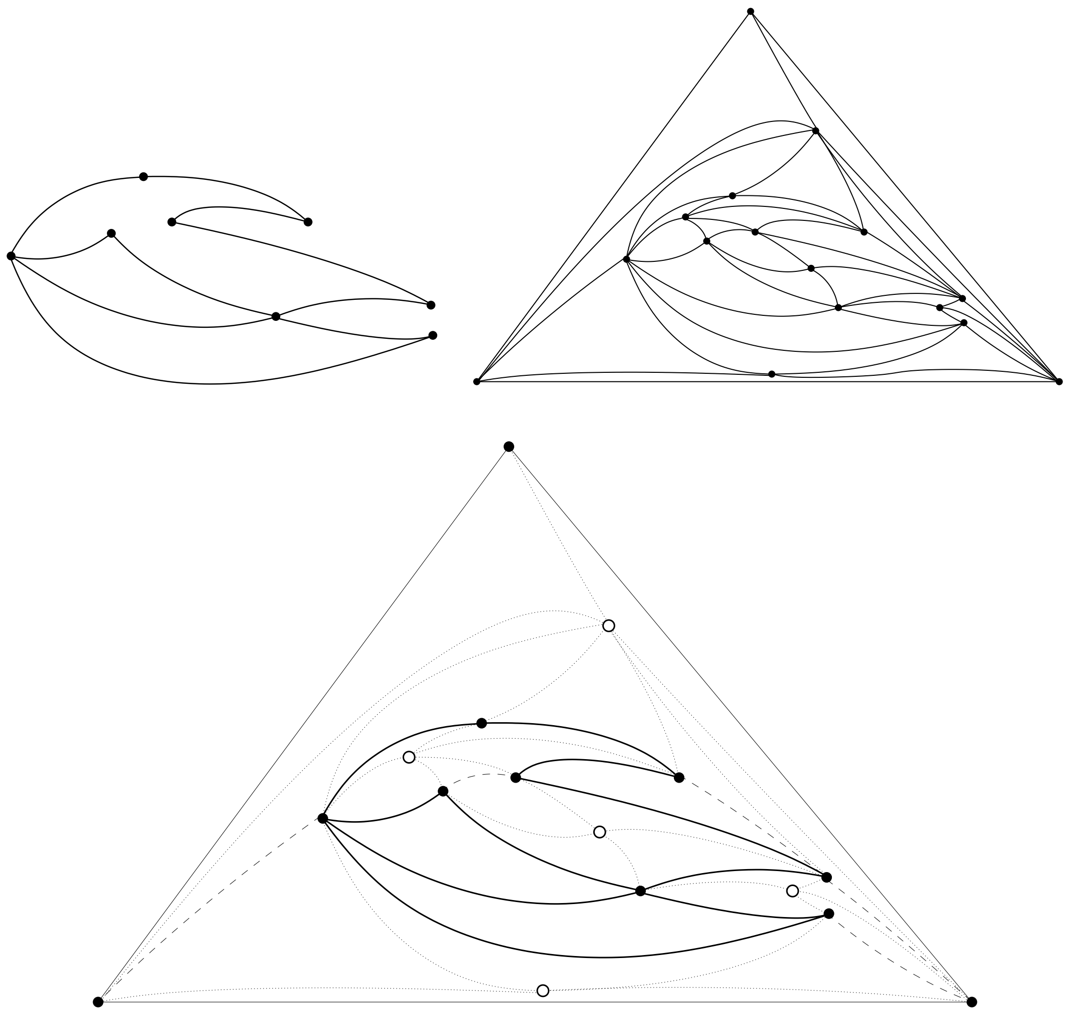

Figure 3.4: Extension of an x-monotone representation of a 1D self-touching configuration to a graph fulfilling the conditions of Lemma 3.1. Top left: Original x-monotone representation. Top right: The resulting graph. Bottom: The extension. In dashed lines, transformation for fulfilling (b); dotted lines and white vertices, transformation for fulfilling (a).

edges are not vertical. This is obvious for the edges added for condition (b). The edges added for condition (c) are not vertical if we place conveniently the vertices in the interior of the non-triangular faces. □

For proving Theorem 3.3, which is used for general self-touching configurations, we need the following variation of Lemma 3.1.

Lemma 3.2. *Any x-monotone representation of a one-dimensional self-touching configuration with an outer face without chords and consisting of two x-monotone chains, can be extended to an x-monotone planar graph such that:*

(a) It is 3-connected.

(b) Every interior vertex, that is, every vertex not incident to the exterior face, has left and right neighbors.

Proof. We can extend the planar embedding to a graph fulfilling (b) as described in the proof of Lemma 3.1. Achieving 3-connectivity is not a problem for the interior vertices since we triangulate the interior of the graph. The graph cannot become disconnected by removing two vertices: it could only be disconnected by removing two boundary vertices, but since we do not have chords this cannot happen. □

3.2.1 Tutte Embedding for Fixed x-Coordinates

Given a graph $G = (V, E)$ fulfilling the conditions of Lemma 3.1, we can apply the Tutte embedding.

The classical Tutte embedding described in Section 0.2 finds the equilibrium position of the interior vertices of G for a fixed convex boundary and an assignment of positive stresses ω_{ij} to the interior edges.

In this chapter we use the following variant of the Tutte embedding: we fix the convex boundary and the x-coordinates of the interior vertices, and we search for an assignment of stresses to the interior edges and the y-coordinates of the interior vertices such that the configuration is in equilibrium.

First of all, we fix the coordinates of our triangular boundary $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$, where $\mathbf{p}_1$ and $\mathbf{p}_2$ are the leftmost and rightmost vertices respectively. For computing the equilibrium stress and the y-coordinates of the $n-3$ interior vertices that fit with the fixed x-coordinates, we apply the approach described in Chrobak et al [14]. We say that a stress is in *x-equilibrium* if it satisfies the equilibrium equations for the x-coordinate. The next lemma shows how to compute, given the x-coordinates, a positive x-equilibrium stress, i. e., an x-equilibrium stress that is positive for all interior edges.

Define the *x-cost* c_{ij} of an edge $\{i, j\}$ to be $|\omega_{ij}(x_i - x_j)|$.

Lemma 3.3. *Given an embedded 3-connected planar graph G of n vertices with a convex outer face $\mathbf{p}_1, \ldots, \mathbf{p}_k$ and no vertical edges, we can compute, in $O(n)$ time, a positive x-equilibrium stress on G such that each x-cost c_{ij} is a positive integer with magnitude $O(n)$.*

This lemma is proven in [14] for the particular case in which the graph has a triangular boundary, and it can be analogously proven for any convex boundary. The sketch of the proof is the following. Let us orient all edges of G from left to right (this is well-defined, since G

contains no vertical edges); this is, we denote by (i, j) an oriented edge, with $x_i < x_j$. Let $\mathbf{p}_1$ and $\mathbf{p}_2$ be the leftmost and rightmost boundary vertices respectively. By our assumptions, for each interior edge (i, j), there exists an x-monotone directed path P_{ij} from $\mathbf{p}_1$ to $\mathbf{p}_2$ containing (i, j). The x-cost on each edge can be seen as a flow from left to right, with the x-equilibrium equation serving the role of flow conservation at each node. However, we do not set any capacity constraint on the edges. The initial flow is 0 on all edges. Then, for each interior edge (i, j) we increase by one unit the flow along the path P_{ij} from $\mathbf{p}_1$ to $\mathbf{p}_2$. Since we maintain the interior x-equilibrium with each "augmentation", this procedure results in an interior-positive stress in equilibrium. The stresses are obtained as $\omega_{ij} = c_{ij}/(x_i - x_j)$. The method described here works in time $O(n^2)$. Chrobak, Goodrich and Tamassia [14] show how to achieve a running time of $O(n)$ by carefully picking the augmenting paths P_{ij}. They also prove that all x-costs in G are integers bounded by $O(n)$.

Hence, from the x-coordinates, we can determine by Lemma 3.3 the interior positive stresses ω_{ij}. Thus, following the notation given in Section 0.2, $\bar{L}$ is known and solving the equilibrium system for the y-coordinate

$$\bar{L} \cdot \mathbf{y} = \mathbf{b}_{\mathbf{y}},$$

we obtain the y-coordinates of the interior points. Thus we have obtained the positions of the straight-line embedding with the imposed x-coordinates.

3.2.2 The δ-Perturbation

Summarizing the previous results we are able to prove Theorem 3.2.

Proof of Theorem 3.2. Use Lemma 3.1 to extend the x-monotone planar embedding to a 3-connected graph with an x-monotone triangular boundary and where every interior vertex has left and right neighbours. Fix the vertices of the triangular boundary within δ-balls centered on the line in which the original self-touching configuration lay. Since we maintain the combinatorial planar embedding, the y-coordinate of the interior vertices will be perturbed less than δ (otherwise they would switch to the exterior of the triangle).

Fix the x-coordinates of all interior vertices at their original position. Apply Lemma 3.3 to compute a positive x-equilibrium stress ω, and solve the equilibrium system for the y-coordinates of the interior points, obtaining the coordinates of the spring embedding. □

Analogously, we can prove Theorem 3.3.

Proof of Theorem 3.3. We extend the planar embedding to a graph fulfilling the conditions of Lemma 3.2. Then we can apply Lemma 3.3 and the rest of the proof is analogous to the proof of Theorem 3.2. □

Given a one-dimensional self-touching configuration, we draw an x-monotone representation, as described at the beginning of this section. Applying Theorem 3.2 and scaling such that all vertices differ at most $\delta/2$ from its original position, we conclude with the following result.

Theorem 3.4. *For every one-dimensional self-touching configuration, and for each $\delta > 0$, there exists a simple δ-perturbation.*

3.3 2-D Self-Touching Configurations

Let $\mathcal{C}$ be any self-touching configuration in $\mathbb{R}^2$. We show Theorem 3.1 by perturbing $\mathcal{C}$ in two stages. The idea is to structure the self-touching configuration into several one-dimensional elongated oval structures that we call *cigars*, so that each cigar contains several overlapping edges. This structuring allows us to perturb the configuration in two stages. In a first stage, we perturb the boundaries of all cigars so that they become convex. In a second stage, we perturb the interior of the cigars using the results of Section 3.2.

We assume that there are no vertical segments in $\mathcal{C}$ by just rotating the configuration a little, i. e., we assume that there are no vertical cigars.

In the following we denote by $B_r(x)$ the open disk of radius r centered at x.

3.3.1 Cigars

For every segment (edge) of the self-touching configuration, we join all overlapping bars into one-dimensional *cigars*, so that each bar belongs exactly to one cigar.

If there is a bar along two or more consecutive segments, then the overlapping bars of these segments are included in the same cigar. If there is no common bar along two parallel consecutive segments, then these two segments belong to different cigars. See Figure 3.5.

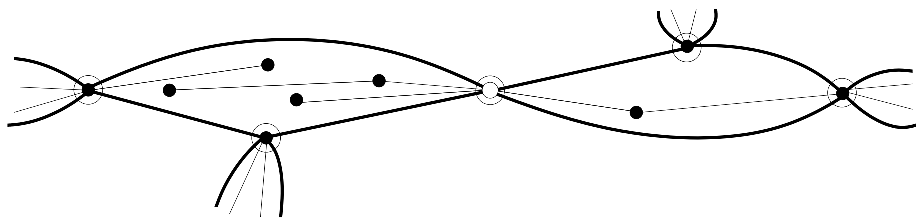

Figure 3.5: Grouping vertices and parallel bars into cigars. Boundary vertices are surrounded by a circle and boundary bars are drawn in thick lines. The vertex in white converges to the left endpoint of one cigar and the right endpoint of the other cigar.

When a bar of the self-touching configuration does not overlap with any other bar, it constitutes a cigar by itself. We call this type of cigars *sticks*.

In each cigar C we distinguish between *interior vertices* and *boundary vertices*.

The *boundary vertices* of a cigar C are the vertices of C which are endpoints of some cigar, which can be C itself or another cigar C'. In other words, vertices belonging to more than one cigar are boundary vertices of the cigars they belong to.

Remark 3.1. *A vertex belonging to t cigars must converge to at least the endpoint of $t-1$ of these cigars. That is, two cigars C_1 and C_2 never share a boundary vertex which does not converge to any of the endpoints of C_1 and C_2.*

Note also that one vertex belongs to more than one cigar if it is incident to bars pointing to several directions, or incident to parallel bars lying along two consecutive parallel segments of $\mathcal{C}$ such that no other bar lies along both segments (in this case the vertex belongs to the boundary of two cigars: it converges to the left endpoint of one cigar and the right endpoint of the other one, see the white point in Figure 3.5).

All the remaining vertices of C are the *interior vertices*. The interior vertices belong only to C and do not converge to an endpoint of C.

We introduce *boundary bars* connecting the boundary vertices so that the boundary of each cigar becomes a flattened self-touching polygon. Obviously we do not introduce those boundary bars already existing in $\mathcal{C}$, otherwise we would have multiple bars.

The boundary bars connecting two boundary vertices converging to the same point have an infinitesimal length. We say that these bars have *zero length*. When there are several converging boundary points, they are connected by zero length boundary bars according to the topological planar embedding of the self-touching configuration, such that the boundary of the cigar has the topology of a planar polygon.

In the case when a cigar that is not a stick has only two boundary vertices, one in each endpoint, we must introduce an artificial boundary vertex, otherwise we would have multiple bars on the boundary of the cigar. See Figure 3.6.

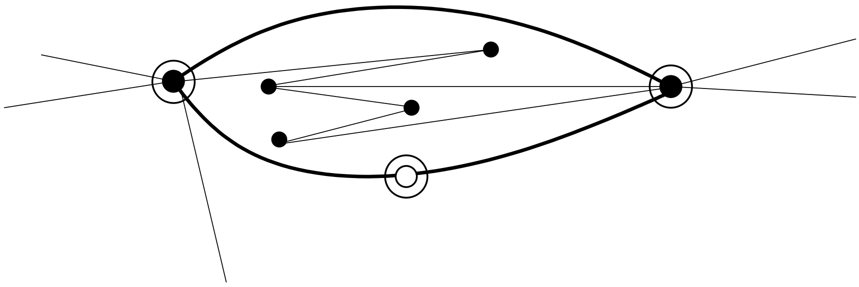

Figure 3.6: Addition of an artificial boundary vertex (in white) for avoiding multiple bars.

The bars adjacent to at least one interior vertex are *interior bars* of the corresponding cigar.

In a cigar, some pairs of boundary vertices which are not connected by a boundary bar may be adjacent. We call *chords* the bars joining two boundary vertices not connected by a boundary bar. Chords are neither interior bars nor boundary bars.

In the case of a stick, its two extreme vertices are its only boundary vertices, and the bar itself is the only boundary bar. (Obviously a stick has an empty interior.)

See in Figure 3.7 an example of how a self-touching configuration is structured into cigars.

3.3.2 First Stage: Perturbing the Boundaries of the Cigars

In this first stage we forget about the interior of the cigars. We remove the interior edges and the interior vertices of all cigars, obtaining a self-touching configuration $\mathcal{C}_B$ made up only of the cigar boundaries. See Figure 3.8.

Our goal now is to obtain a δ-perturbation of $\mathcal{C}_B$ where the boundaries of all cigars are convex polygons. We do it in several steps. First we convexify the boundaries of what we call *merged cigars*, without working out the self-touching incidences. Later we perturb the vertices converging to the same point while maintaining the convexity of the cigar boundaries.

The boundary vertices of a cigar are classified into endpoint vertices, lower boundary vertices and upper boundary vertices. The endpoint vertices are those converging to the endpoints of the cigar. The lower and upper boundary vertices are not endpoint vertices.

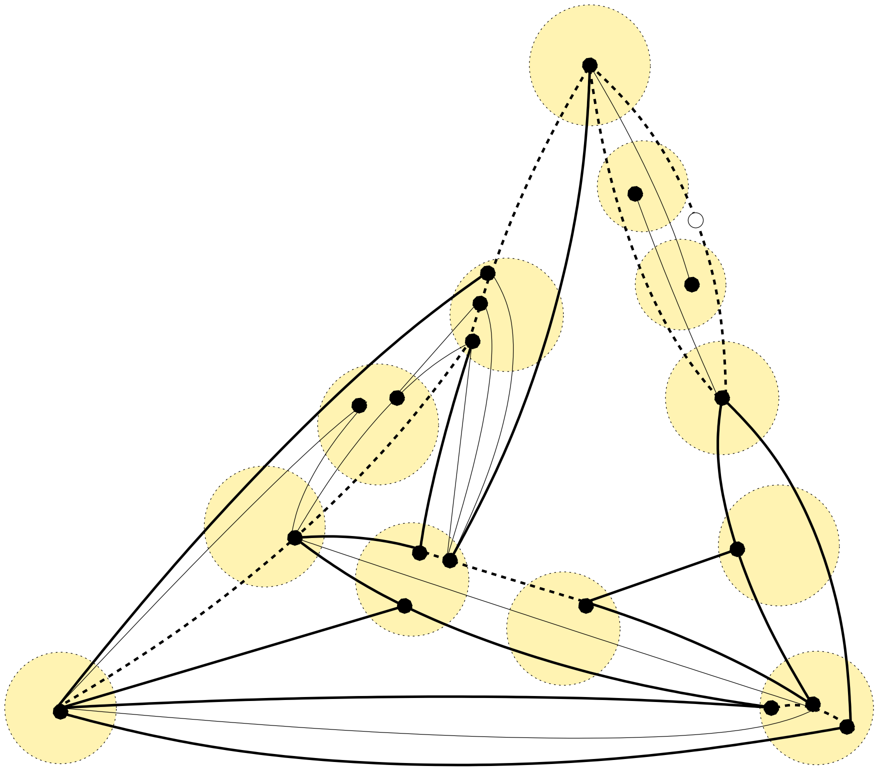

Figure 3.7: A self-touching configuration $\mathcal{C}$ and its cigars. Boundary bars are drawn in thick lines, and those boundary bars introduced artificially are drawn with dashed lines. Boundary vertices introduced artificially are drawn in white. Bars between vertices converging to the same position, that is, between vertices inside the same dotted disk, are zero length. The thin continuous lines represent the interior bars. Bars are drawn as curves instead of straight lines.

Merged Cigars

In $\mathcal{C}_B$, we merge vertices infinitesimally close to each other as follows. See Figure 3.9.

Vertices converging to the same point are merged into a unique vertex, also when they belong to boundaries of different cigars. But we do not merge upper boundary vertices with lower boundary vertices of the same cigar.

In the case when an endpoint x of a cigar C_x converges to the interior of a boundary bar of another cigar C_y, we add a new vertex y, touching x, to the boundary of C_y, and we merge x and y. Let y_p and y_s be respectively the preceding and the succeeding vertex to y on the boundary of C_y. The new bars $\{y, y_p\}$ and $\{y, y_s\}$ become boundary bars, and the existing boundary bar $\{y_p, y_s\}$ becomes a chord.

Note that in the case where C_y is a stick with endpoints y_p and y_s, it becomes a cigar with a triangular boundary of vertices y, y_p, y_s (it is no longer a stick).

We denote by $\mathcal{C}'_B$ the resulting self-touching configuration.

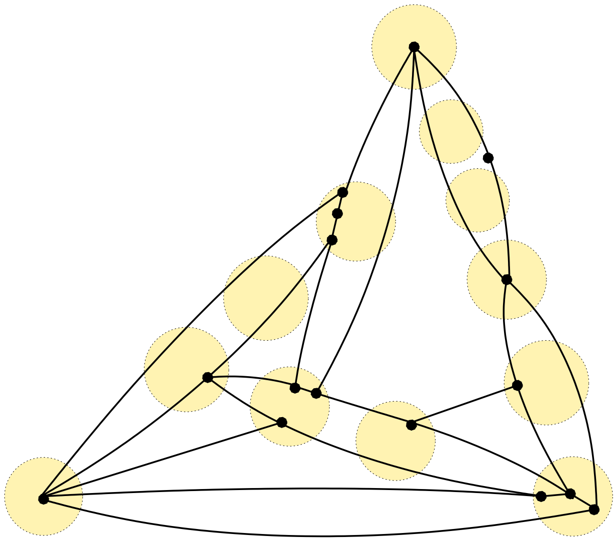

Figure 3.8: The cigar boundaries $\mathcal{C}_\mathcal{B}$ of the self-touching configuration in Figure 3.7.

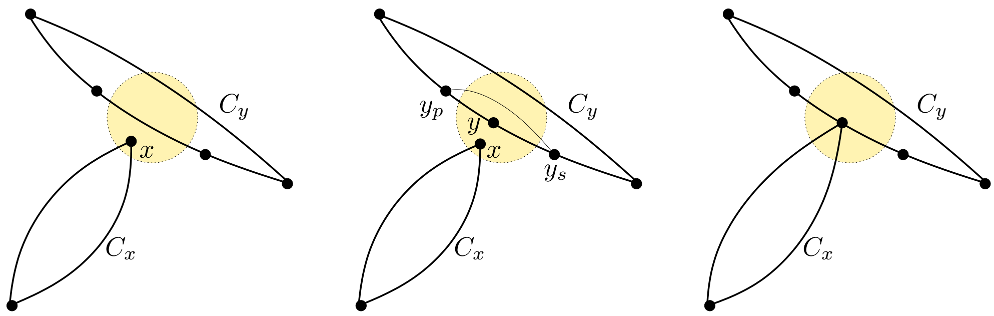

Figure 3.9: Merging an endpoint x of a cigar C_x incident to the interior of a boundary bar of another cigar C_y.

Convexifying the Merged Cigars

We want to obtain a $\delta/3$-perturbation of $\mathcal{C}'_B$ where the boundaries of all cigars are strictly convex polygons (except for the sticks). We perturb now the vertices only within disks of radius $\delta/3$ since later we move some vertices again, so we ensure this way that later we will still have place inside the δ-disks.

Since there are no vertical bars, for each cigar it makes sense to speak of "above" or "below". We say that a cigar A is *below* a cigar B, and we write $A \prec B$, if an endpoint of A coincides with a lower boundary vertex of B, or if an upper boundary vertex of A coincides with an endpoint of B. If $A \prec B$, we can also say that B is *above* A. This is illustrated in Figure 3.11.

In case two cigars A and B have an endpoint in common, A and B are not compared by the relation $\prec$.

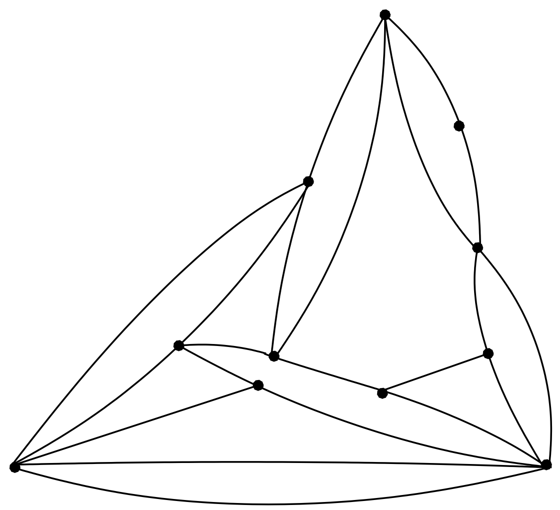

Figure 3.10: The cigar boundaries of Figure 3.8, now merged and constituting $\mathcal{C}'_B$.

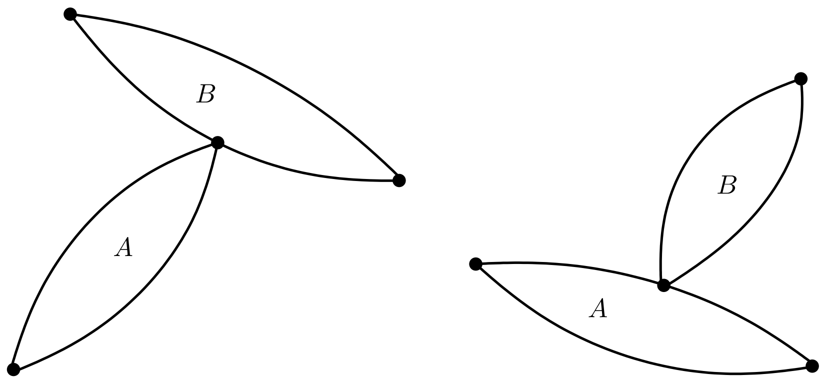

Figure 3.11: Cases for a cigar A below a cigar B, $A \prec B$.

Lemma 3.4. *The relation $\prec$ is acyclic.*

Proof. Suppose it is not true. Then let P be the smallest directed cycle. Suppose in P we have $A \prec B \prec C$, for given cigars A, B and C. Consider a straight arc from the common vertex of A and B to the common vertex of B and C. We draw such a straight arc for all the comparative relations in P. Then, on the self-touching configuration, P is represented by a directed simple polygon. The vertices of the polygon correspond to vertices of $\mathcal{C}'_B$, and the arcs are along original edges of $\mathcal{C}'_B$ (not necessarily along the entire edge). P is simple, i. e., it does not have self-intersections: its arcs P do not cross since the underlying self-touching configuration is planar, and the vertices of P are all different since we are taking the minimum cycle.

The directed polygon P has two possible orientations: counterclockwise or clockwise. Suppose it is oriented counterclockwise. Consider the leftmost vertex l of P.

Analogously to Lemma 2.2, given a simple polygon, consider the directed walk from a bottommost vertex to a topmost vertex along the left boundary chain. If l is a leftmost vertex on this walk with a non-vertical ingoing edge, then the edge leaving l must lie in the lower half-plane defined by the edge entering l. Let A and B the cigars meeting at l, such that the edge entering l lies along the cigar A and the edge leaving l lies along the cigar B. Then l must be either an endpoint of A and an upper boundary vertex of B, or a lower boundary vertex of A and an endpoint of B. Hence it is impossible that $A \prec B$, thus the cycle P cannot exist.

The proof is analogous for P oriented clockwise, by considering the rightmost point of P. □

We convexify the cigar boundaries in two steps. First we take the cigars of $\mathcal{C}'_B$ from bottom to top according to the relation $\prec$, so that if $A \prec B$ then A is taken before B, and we convexify their upper boundaries one by one. Second, we take the cigars of $\mathcal{C}'_B$ in the inverse order, from top to bottom, so that if $A \prec B$ then B is taken before A, and we convexify their lower boundaries one by one.

Let us then convexify the upper boundaries one by one in the mentioned order, from bottom to top. For a cigar C, the idea is to move each vertex v of C vertically upwards a certain amount Δy_v^+, which is defined inductively. The amount Δy_v^+ is positive for all upper boundary vertices of C and endpoint vertices of C which are upper boundary vertices of some other cigars, and zero for the remaining vertices, that is, lower boundary vertices and endpoint vertices which are not upper boundary vertices of any other cigar.

Let i and j be the endpoint vertices of C. If i is an upper boundary vertex of some other cigar C', then we have $C' \prec C$, hence $\Delta y_i^+ > 0$ has been already defined. Otherwise, we set $\Delta y_i^+ = 0$. Analogously, $\Delta y_j^+ > 0$ has been already defined as an upper boundary vertex of a preceding cigar, or it is set to $\Delta y_j^+ = 0$. Once Δy_i^+ and Δy_j^+ are defined for both endpoint vertices i and j, we define Δy_k^+ for any upper boundary vertex k of C as

$$\Delta y_k^+ := \frac{x}{l}\Delta y_j^+ + \frac{l-x}{l}\Delta y_i^+ + x(l-x),$$

where l is the distance from i to j, and x is the distance from i to k. The first two summands translate the vertical movement of the endpoints i and j to the upper boundary vertex k. The term $x(l-x)$ ensures that the upper boundary is strictly convex after moving all the upper boundary vertices of C. See Figure 3.12.

When we have visited all cigars from bottom to top, Δy_v^+ is defined for all vertices of $\mathcal{C}'_B$. Then we visit all cigars one by one from top to bottom. The idea is to move each vertex v vertically downwards a certain amount Δy_v^- defined inductively. The amount Δy_v^- is negative for all lower boundary vertices, and zero for the remaining vertices, that is, for all upper boundary vertices and endpoint vertices which are not lower boundary vertices of any cigar.

Suppose Δy_i^- and Δy_j^- are already defined for the endpoint vertices i and j of a cigar C. Then, for any lower boundary vertex k of C, Δy_k^- is defined as

$$\Delta y_k^- := \frac{x}{l}\Delta y_j^- + \frac{l-x}{l}\Delta y_i^- - x(l-x),$$

where l is the distance from i to j, and x is the distance from i to k. Analogously to the previous case, the first two summands translate the vertical movement of the endpoints i and j to k, while $-x(l-x)$ ensures the strict convexity of the lower boundary.

Finally, each vertex v of $\mathcal{C}'_B$ with coordinates (x_v, y_v) is moved vertically to $(x_v, y_v + \varepsilon\,(\Delta y_v^+ + \Delta y_v^-))$. We choose the parameter ε small enough such that v is moved within a disk of radius $\delta/3$ centered at (x_v, y_v), i. e, such that $|\varepsilon\,(\Delta y_v^+ + \Delta y_v^-)\,| \leq \delta/3$. See Figure 3.13.

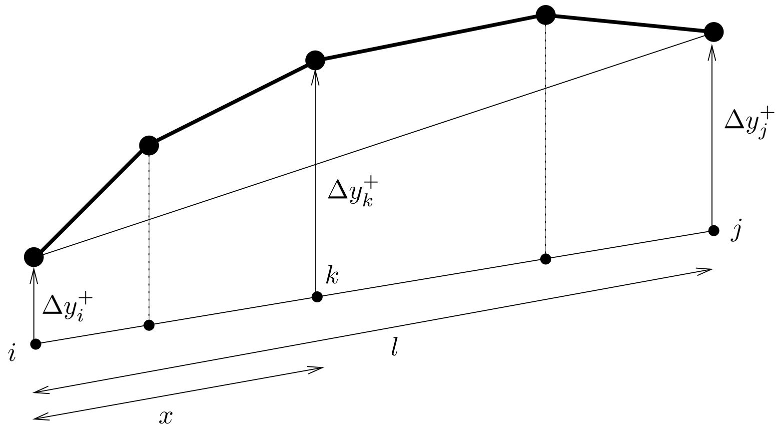

Figure 3.12: This is how the upper boundary of a cigar would look (thick lines) if we would move each of its vertices $v = (x_v, y_v)$ to $(x_v, y_v + \varepsilon\Delta y_v^+)$. In this case, both endpoints of the cigar are also upper boundary vertices of some other cigar, since they are moved vertically upwards a positive amount. Small points represent the original positions of the vertices.

Dealing with the Touching Vertices of the Boundaries

Once the merged cigars have a strictly convex boundary we have to first "unmerge them" again, and second, draw the δ-perturbation of the self-touching configuration $\mathcal{C}_B$. This can be done independently for each vertex v of $\mathcal{C}'_B$.

Take a vertex v of $\mathcal{C}'_B$. This vertex may include many merged vertices from several cigars.

Step 1. Separating Cigars without Vertices in Common. As a first step we "unmerge the cigars", that is, we separate cigars and sets of cigars with no vertices in common. We want to spread the vertex v into possibly several vertices such that each cigar previously containing v in $\mathcal{C}'_B$ contains exactly one vertex that converged to v. This spreading is performed again within the interior of a disk of radius $\delta/3$ centered at v.

By Remark 3.1, at most one vertex converging to v is not an endpoint of a cigar. Let u^0 be such a vertex, if it exists. All other vertices converging to v, denote them by $u_1, \ldots, u_t$, are endpoints of cigars.

The strategy is as follows. In case it exists, u^0 remains fixed. For perturbing the rest of the vertices $u_1, \ldots, u_t$ we proceed as follows. Each vertex u_i, $1 \leq i \leq t$, can be an endpoint of several cigars $C_i^1, \ldots, C_i^{s_i}$. The goal is to spread the vertices u_i while maintaining the strict convexity of the cigars $C_i^1, \ldots, C_i^{s_i}$.

For each vertex u_i, and for each j, $1 \leq j \leq s_i$, there is an area where u_i can be placed such that C_i^j remains convex. We denote by $\mathcal{A}(C_i^j, u_i)$ the interior of this area.

In the general case when both upper and lower boundaries of C_i^j contain two or more vertices, $\mathcal{A}(C_i^j, u_i)$ is the interior of the area between the lines kl, mn and ln, containing u_i, where l and k (resp. n and m) are respectively the first and the second vertices after u_i on the upper (resp. lower) boundary of C_i^j (see Figure 3.14). Note that if the lines kl and mn intersect in the same halfplane defined by ln that contains u_i, then $\mathcal{A}(C_i^j, u_i)$ is bounded by the triangle formed by the lines kl, mn, ln.

In the case when the upper (resp. lower) boundary is just a bar joining the two endpoints,

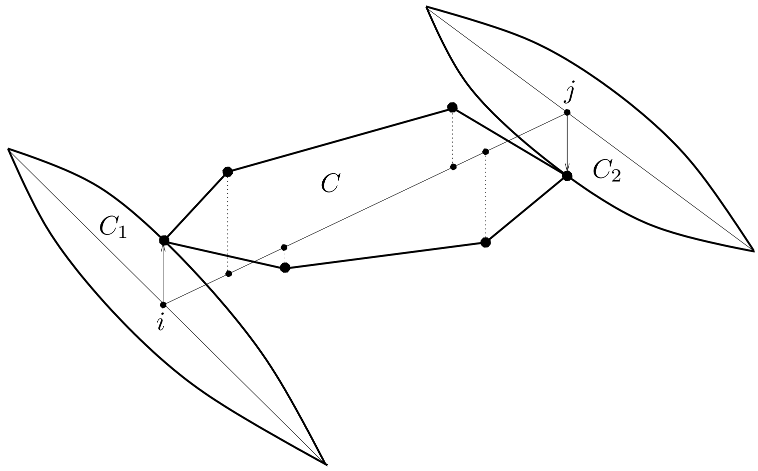

Figure 3.13: A cigar C, strictly convex after moving vertically its boundary vertices. The endpoint vertex i is an upper boundary vertex of another cigar C_1, such that $C_1 \prec C$, hence $\Delta y_i^+ > 0$. The other endpoint vertex j is a lower boundary vertex of another cigar C_2, such that $C \prec C_2$, hence $\Delta y_j^- < 0$.

then $\mathcal{A}(C_i^j, u_i)$ is defined as the open halfplane delimited by mn (resp. kl) and containing u_i. This is also necessary even when the cigar has a triangular boundary, which is always convex, since, in order to maintain the combinatorial planar embedding, u_i cannot switch to the other halfplane. In the case when C_i^j is a stick, $\mathcal{A}(C_i^j, u_i)$ is the whole plane, since a stick is always convex.

Note that since C_i^j is strictly convex, u_i belongs to $\mathcal{A}(C_i^j, u_i)$.

For perturbing u_i within $\delta/3$-disks while maintaining the convexity of all the cigars $C_i^1, \ldots, C_i^{s_i}$ we must consider the intersection

$$\mathcal{R}_{u_i} = \left(\bigcap_{j=1}^{s_i} \mathcal{A}(C_i^j, u_i) \right) \cap B_{\delta/3}(u_i)\,.$$

This intersection is non-empty since the original position of u_i, v, is interior to each $\mathcal{A}(C_i^j, u_i)$ and interior to $B_{\delta/3}(u_i)$. We obtain a non-empty open region $\mathcal{R}_{u_i}$ where u_i can be drawn, for each vertex u_i, $1 \le i \le t$. See Figure 3.15 for an example.

Now we must spread all vertices u_i, $1 \le i \le t$, each within $\mathcal{R}_{u_i}$, ensuring that no crossings between cigars are introduced. Let the *sector* S_i of a vertex u_i be the smallest circular sector of the disk $B_{\delta/3}(v)$ containing all the cigars $C_i^1, \ldots, C_i^{s_i}$. We say that a vertex u_i is *nested* within a vertex u_j if S_i is contained in S_j. See Figure 3.16 (Left). Note that, if we have $S_i = S_j$, then one of the two vertices, say u_j, is the endpoint of two sticks. In this case, we say that u_i is nested within u_j. Note also that, for each u_i, its corresponding sector S_i intersects $\mathcal{R}_{u_i}$. We perturb each vertex u_i within $\mathcal{R}_{u_i}$ and into the interior of its sector S_i. If u_i is nested within another vertex u_j, then u_i must be perturbed further away from its initial position than u_j, to ensure that we leave sufficient space to perturb u_j. This is no problem since $\mathcal{R}_{u_i}$ and $\mathcal{R}_{u_j}$ have a non-empty open intersection containing v (because they are both open regions containing v).

At the end, $u^0, u_1, \ldots, u_t$ are not touching anymore. See Figure 3.16 (Right).

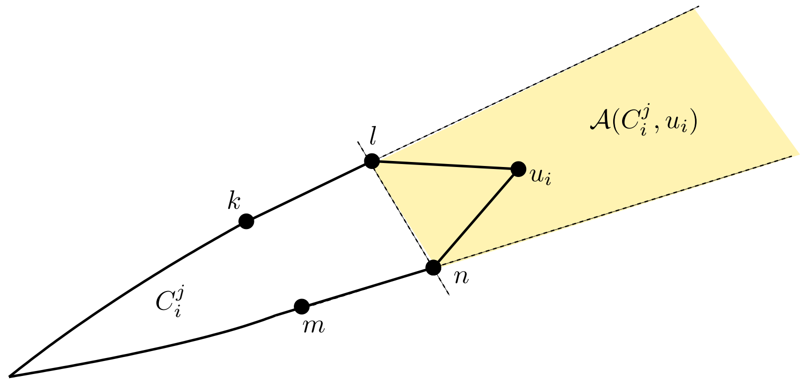

Figure 3.14: A cigar C_i^j with endpoint u_i. The shaded area $\mathcal{A}(C_i^j, u_i)$ is the region where u_i can be moved while maintaining the convexity of C_i^j.

Step 2. Perturbing the Remaining Converging Vertices. Now we must perturb the remaining vertices of $\mathcal{C}_B$ converging to u^0 and to each $u_1 \ldots, u_t$, independently. Let u be one of the vertices $u^0, u_1, \ldots, u_t$. In this step, we want to spread all vertices converging to u.

Remark 3.2. *The set of cigars is connected, otherwise they would have been already separated in Step 1.*

After the perturbation, all adjacencies must be preserved and each cigar must be a convex closed polygon. Strict convexity is needed at those vertices incident to at least a boundary bar which did not have zero length in $\mathcal{C}_\mathcal{B}$. This is to avoid that chords overlap with the boundary bars and between themselves. Those boundary vertices which in $\mathcal{C}_\mathcal{B}$ were incident to two zero length bars need not be strictly convex. See Figure 3.17. Note that the zero length boundary bars between vertices previously touching become short bars after the perturbation.

In this step we allow again every vertex to move within the interior a disk of radius $\delta/3$ centered at its current position. As before we play with the fact that, for an endpoint x of a convex cigar C which converges to u, there is an open area $\mathcal{A}(C, u)$ where x can be drawn maintaining the convexity of C. Let $C_1, \ldots, C_{s_u}$ be the set of cigars with an endpoint converging to u. For perturbing x within $B_{\delta/3}(u)$ while maintaining the convexity of all the cigars with endpoint x, we consider the intersection of all $\mathcal{A}(C_j, u)$, $j = 1, \ldots, s_u$.

$$\mathcal{R}_u = \left(\bigcap_{j=1}^{s_u} \mathcal{A}(C_j, u) \right) \cap B_{\delta/3}(u)\,.$$

Since u is interior to each of these areas, the open region $\mathcal{R}_u$ is nonempty and there is always space enough to perturb new vertices inside $\mathcal{R}_u$.

The vertices converging to u are perturbed successively according to a hierarchy which is defined later. Hence, when we look at some vertices and think how to perturb them, some other vertices are already drawn in their final positions.

We give some rules for perturbing the vertices converging to u, although it can be done in many ways. First we describe a procedure which is used later during the perturbation. Given

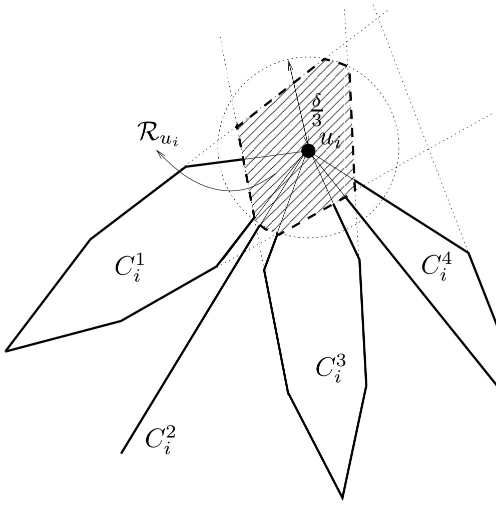

Figure 3.15: In the example, u_i is the endpoint of four cigars $C_i^1, \ldots, C_i^4$. The shaded area represents the open region $\mathcal{R}_{u_i}$ where u_i can be moved while maintaining the convexity of $C_i^1, \ldots, C_i^4$.

two cigars C_a and C_b that share at least one vertex already drawn in its final position, this procedure describes how to perturb the remaining common endpoints.

Procedure A. Perturbing the common endpoints of two cigars C_a and C_b sharing a vertex already drawn in its final position. Assume that the cigars C_a and C_b share a vertex y already drawn in its final position. Let $y, y_1, \ldots, y_s$ be the endpoint vertices shared by C_a and C_b. We choose a ray R starting at y and lying between C_a and C_b. We place $y_1, \ldots, y_s$ on R (in this order, so that y_s is the furthest vertex from y), within $\mathcal{R}_u$, redrawing the end of the cigars C_a and C_b. See Figure 3.18. Note that both cigars are strictly convex at y_s. Note that the vertex y_s can be also an endpoint vertex of some other cigars $D_1, \ldots, D_r$ lying between C_a and C_b. Note also that, since $y_1, \ldots, y_s$ are redrawn within $\mathcal{R}_u$, the convexity of $C_a, C_b, D_1, \ldots, D_r$ is maintained.

For spreading all vertices converging to u the strategy is as follows. We define a hierarchy that gives the order in which we draw the vertices converging to u and their corresponding cigars, whose ends must be redrawn. We classify the vertices converging to u into *levels*. We draw the vertices level by level: First we perturb the vertices belonging to the *starting level* or level 0. Then we perturb the vertices of the next level, and so on.

Definition of the hierarchy. Let x be a vertex converging to u. As in Step 1, let the *sector* S_x of x be the smallest circular sector of the disk $B_{\delta/3}(u)$ containing all the cigars with endpoint x. If the angle of this sector is greater than π (this corresponds to the case when x is the endpoint of several cigars not contained in a halfplane), then we define S_x to be the whole disk $B_{\delta/3}(u)$. We say that a vertex x is *nested* within a vertex y if S_x is contained in S_y.

The vertices belonging to the *starting level*, or level 0, are those vertices converging to u which are not nested within any other vertex. Note that when $u = u^0$ is a vertex of the upper

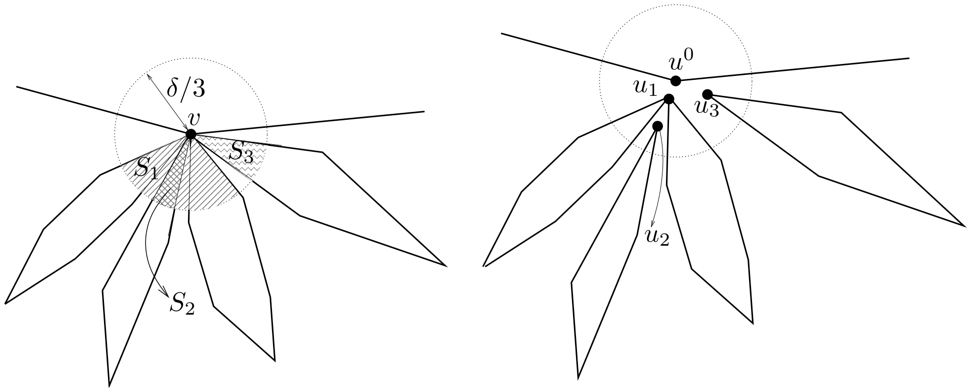

Figure 3.16: Left: The sectors S_1, S_2, S_3 of the endpoint vertices converging to v. S_1 is the striped area, S_2 is the checked area, and S_3 is the area in zigzag. Since $S_2 \subset S_1$, the vertex u_2 is nested within the vertex u_1. Right: Spreading the vertices u^0, u_1, u_2, u_3.

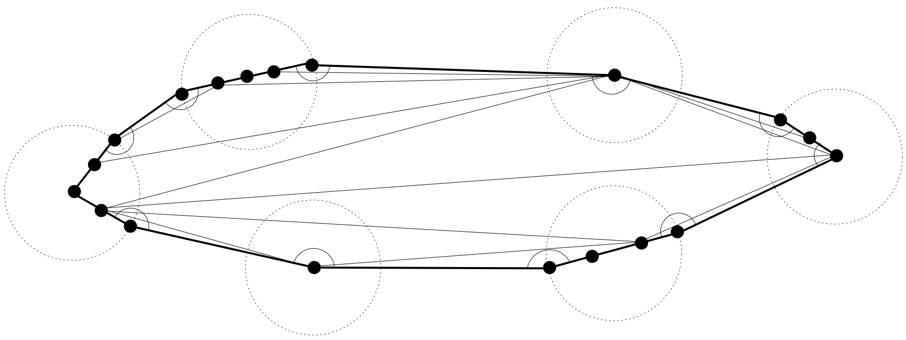

Figure 3.17: Schematic drawing of a cigar after Step 2. Dotted circles enclose vertices which converge to the same geometric position in $\mathcal{C}_{\mathcal{B}}$. Chords are drawn in thin lines, and marked angles must be strictly convex. Actually the cigar is much flatter and dotted circles are much smaller.

or lower boundary of some cigar A, the vertices of A converging to u^0 belong to the starting level, since they cannot be nested within any other vertex.

The level i contains the vertices which are *nested* i times: Suppose the level $i-1$ is already defined. Then, at level i we have the vertices nested within the vertices of the $(i-1)$th level.

There are vertices which belong only to one cigar, and they can be treated easily at the very end, after all vertices have been drawn.

Perturbing the vertices of the starting level.

Case 1: $u = u^0$ is a vertex of the upper or lower boundary of a cigar A. Let $x_1, \ldots, x_r$ be the vertices of A converging to u^0. They all belong to the starting level. By definition of boundary vertex, $x_1, \ldots, x_r$ are also endpoint vertices of other cigars $C_1, \ldots, C_p$ (assume they appear in this order). The vertices shared by two consecutive cigars C_j and C_{j+1}, $1 \leq j \leq p-1$, also belong to the starting level.

Cut the cigar A at u^0 by a segment L, sufficiently close to u^0, and place the vertices

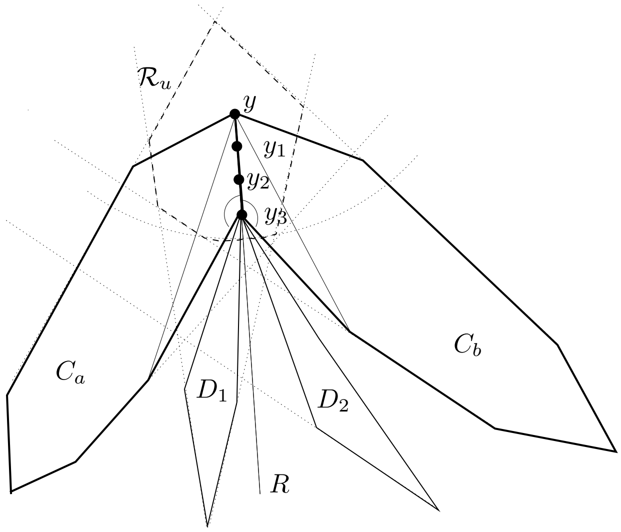

Figure 3.18: Procedure A. Placing on a common line R the endpoint vertices y_1, y_2, y_3 common to two cigars C_a and C_b. The original shape of C_a and C_b is represented by thin lines. The angle between C_a and C_b at y_3 is less than π. The marked angles must be strictly convex. The boundary of the region $\mathcal{R}_u$ is drawn in dashed lines. In the picture, $s = 3$ and $r = 2$.

$x_1, \dots, x_r$ on $L \cap \mathcal{R}_{u^0}$, such that one of these vertices lies on each of the two extremes of L. See Figure 3.19. After this perturbation, the convexity of the cigars $C_1, \dots, C_p$ is maintained. The convexity of A is also maintained because if we cut a vertex of a convex polygon by a segment, the polygon remains convex.

Suppose there are vertices $y_1, \dots, y_s$ at the starting level that do not belong to A but they are common to two consecutive cigars C_j and C_{j+1} that have some endpoint vertices on L, $1 \leq j \leq p-1$. If C_j and C_{j+1} meet at an endpoint vertex lying on L, then $y_1, \dots, y_s$ are perturbed as described in Procedure A.

If, on the contrary, C_j and C_{j+1} do not share any endpoint vertex on L, then let x^j be the vertex of $C_j \cap L$ nearest to C_{j+1}, and let x^{j+1} be the vertex of $C_{j+1} \cap L$ nearest to C_j. We choose a ray R starting at a point between x^j and x^{j+1} on L, and lying between C_j and C_{j+1}, and we place $y_1, \dots, y_s$ on R (but not on L) within $\mathcal{R}_{u^0}$. We redraw the end of the cigars C_j and C_{j+1} and connect the nearest vertex to L, say y_1, to x^j and x^{j+1} by two short edges. The furthest vertex from y, say y_s, can be also an endpoint vertex of some other cigars $D_1, \dots, D_r$ lying between C_j and C_{j+1}. The required convexity for the cigars is maintained and C_j and C_{j+1} are strictly convex at y_s. See Figure 3.20.

Case 2: u is one of the vertices $u_1, \dots, u_t$ from Step 1, and a vertex converging to u is the endpoint of several cigars not contained in a halfplane. Let x be the vertex converging to u which is the endpoint of several cigars $C_1, \dots, C_p$ not contained in a halfplane, $p \geq 3$. Then we leave x fixed. Note that, by planarity, there exist at most one such a vertex, and obviously it belongs to the starting level. Hence, x is the only vertex of the starting level, since S_x is defined to be the whole disk $B_{\delta/3}(u)$.

Case 3: u is one of the vertices $u_1, \dots, u_t$ from Step 1, and all cigars with endpoint con-

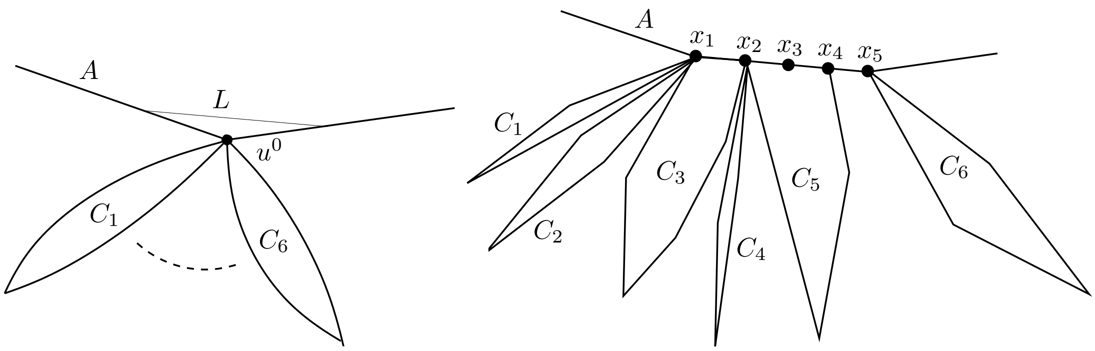

Figure 3.19: Perturbing the touching vertices x_1, x_2, x_3, x_4, x_5, which are non-endpoint vertices of the cigar A and also endpoint vertices of the cigars $C_1, \ldots, C_6$.

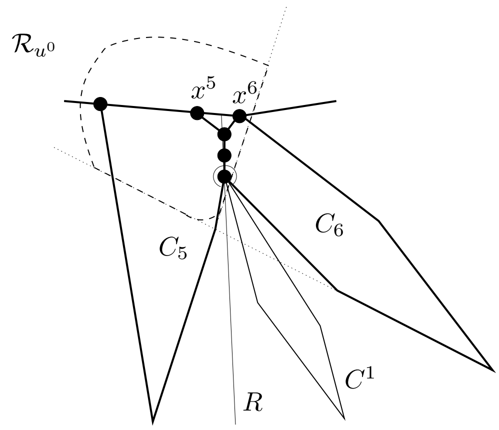

Figure 3.20: Perturbing the vertices shared by C_5 and C_6 of Figure 3.19. Here, $x^5 = x_4$ and $x^6 = x_5$.

verging to u are contained in a halfplane. In this case we may have at the starting level 0 a sequence of cigars $C^1, \ldots, C^q$, $q \geq 3$, (assume they are ordered counterclockwise around u) with endpoints converging to u, such that C^j has some vertices in common with the next cigar and some other different vertices in common with the preceding cigar, $2 \leq j \leq q-1$. See Figure 3.21. The cigar C^q can share some vertices with C^1 or not. In the affirmative case, these vertices are different from the non-empty set of vertices shared by C^1 and C^2, and the non-empty set of vertices shared by C^{q-1} and C^q.

We want to perturb now the endpoints of $C^1, \ldots, C^q$. Let v_j be the vertex topologically nearest to u, shared by two consecutive cigars C^j and C^{j+1} (C^q and C^1 are also considered if they share some vertex). There can be other cigars $D_1, \ldots, D_r$ between C^j and C^{j+1} with endpoint v_j. We redraw the end of each cigar C^j as follows. Consider, for each C^j, a point l_j and a point r_j which will belong to the left and right boundary of C^j respectively (left and right when looking towards the center), such that $r_j = l_{j+1}$ lies within $\mathcal{R}_u$. We draw a segment from l_j to r_j. We place v_j on $r_j = l_{j+1}$, and the rest of vertices shared by two consecutive

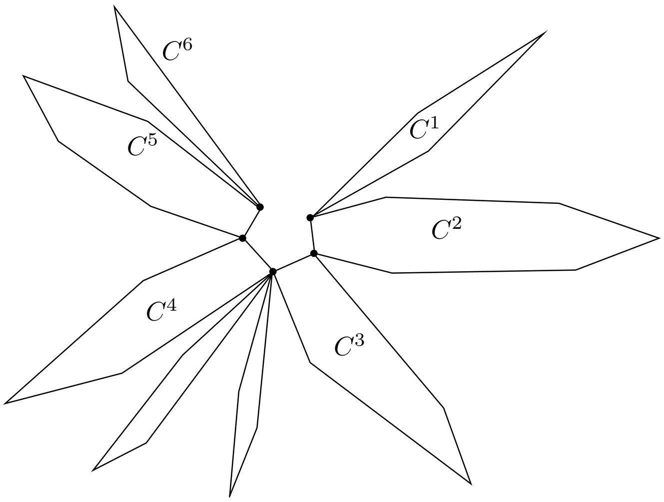

Figure 3.21: Starting level: case when no vertex converging to u is the endpoint of several cigars not contained in a halfplane. In this example, $q = 6$.

cigars lying between C^j and C^{j+1}, that is, vertices between C^j and D_1, between D_k and D_{k+1}, or between D_r and C^{j+1}, are perturbed as described in Procedure A. See Figure 3.22.

If such a sequence of cigars $C^1, \ldots, C^q$ does not exist, then there are three possibilities. First, there are at least three cigars, contained in a halfplane, with a common endpoint x. In this case, x is the unique vertex of level 0, since by planarity x is the unique vertex contained in all these cigars. Then we leave x fixed, as in Case 2. Second, the vertex x topologically closest to u (if there are more than one such vertices, choose one of them), is the endpoint of only two cigars C_a and C_b. Then we leave x fixed, and perturb the other common vertices to C_a and C_b, if they exist, as in Procedure A. Third, a single cigar C has an end converging to u. We place the endpoint vertices of C converging to u within $\mathcal{R}_u$, ensuring that C is strictly convex on those vertices incident to a boundary bar which had non-zero length in $\mathcal{C}_B$.

Perturbing the vertices of level i. Assume the vertices of level $i-1$ are already perturbed, and we are going to perturb the vertices of level i. All vertices of level i are nested between two cigars sharing vertices of level $i-1$ and forming an angle less than π. For any two such cigars, the vertices in between are perturbed independently.

Let C_a and C_b be two cigars forming an angle less than π which are consecutive in level $i-1$ and share some vertices of level $i-1$. We describe how to perturb the vertices at level i between C_a and C_b.

If C_a and C_b share some vertices at level i, we perturb them as described in Procedure A.

If C_a and C_b have no vertex in common at level i, let x be the vertex belonging to both C_a and C_b which has been perturbed furthest from u in level $i-1$. Let $C^1, \ldots, C^q$ be a sequence of cigars between C_a and C_b (assume they are ordered counterclockwise around u) such that C^j and C^{j+1} share some vertices which are different from the set of vertices shared by C^j and C^{j-1}, $2 \le j \le q-1$. Analogously, we assume C_a and C^1 (resp. C^q and C_b) share some vertices, which are different from the set of vertices shared by C^1 and C^2 (resp. C^q and C^{q-1}).

Note that there can exist a j_0, $1 \le j_0 \le q-1$, such that C^{j_0} and C^{j_0+1} are disjoint. (It

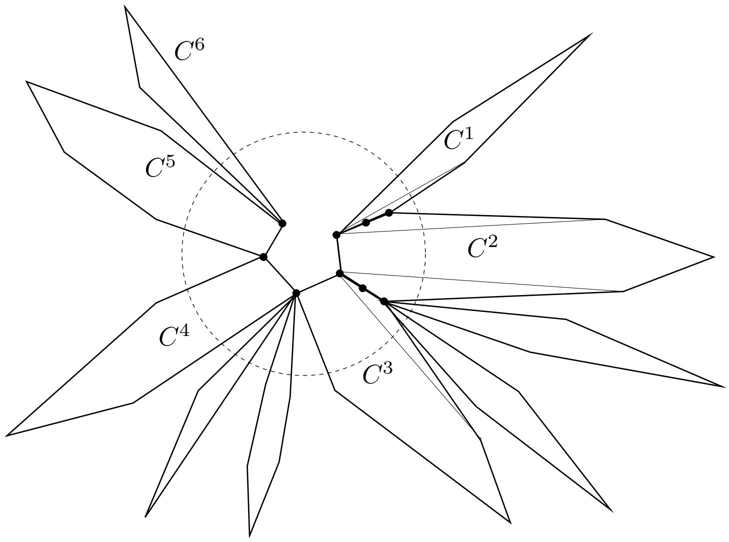

Figure 3.22: How a perturbation at level 0 of the example in Figure 3.21 can look like. This is a schematic drawing. The dashed line represents $B_{\delta/3}(u)$ and separates the vertices converging to u from the vertices converging to other positions. This circle is actually much smaller and cigars are much flatter.

can also happen that C_a and C^1 are disjoint, or C^q and C_b.) In this case, there are no cigars between C^{j_0} and C^{j_0+1}, neither in this level nor in later levels. Such a j_0 can only be found once: if there would exist another j_0', $j_0 < j_0'$, such that $C^{j_0'}$ and $C^{j_0'+1}$ were disjoint, then the set of cigars between C^{j_0+1} and $C^{j_0'}$ would not share any vertex with the rest of cigars converging to u. Hence it would have been already perturbed away in Step 1.

For perturbing the endpoints of $C^1, \ldots, C^q$ we proceed similarly as before. Let v_j be the nearest vertex to x shared by C^j and C^{j+1}, $1 \leq j \leq q-1$. There can be cigars $D_1, \ldots, D_r$ between C^j and C^{j+1} with endpoint v_j. We redraw the end of each cigar C^j as follows. Consider, for each C^j, a point l_j and a point r_j which will belong to the left and right boundary of C^j respectively (left and right when looking towards the center), such that $r_j = l_{j+1}$ lies within $\mathcal{R}_u$. We draw a segment from l_j to r_j. We place the vertices converging to u that belong only to C^j in the interior of this segment, although they actually belong to a posterior level. We place v_j on $r_j = l_{j+1}$. The remaining vertices of level i shared by two consecutive cigars C^j and C^{j+1}, are perturbed as described in Procedure A. Note that we find these vertices at level i only when there are no other cigars between C^j and C^{j+1} with endpoint v_j.

Let v_a be the nearest vertex to x shared by C_a and C^1. The vertex v_a cannot be placed on the current boundary of C_a, since any remaining vertices $y_1, \ldots, y_s$ shared by C_a and C^1 will be perturbed as described in Procedure A, and then C_a would have a non-convex boundary. Note that a solution is not to place $y_1, \ldots, y_s$ also on the boundary of C_a, because then a chord joining the next vertex z of the boundary of C_a, not converging to u, and any of the vertices $v_a, y_1, \ldots, y_{s-1}$, would overlap with the boundary bar $\{z, y_s\}$ (The boundary of C_a must be then strictly convex at y_s). We place v_a on a ray R starting at x and lying between C_a and C_1. The rest of the vertices of level i shared by C_a and C_1 are perturbed as in Procedure A, on the same ray R. Note that if there are cigars between C_a and C_1 with endpoint v_a, then any vertex between C_a and C_1 different from v_a must belong to a posterior level. We perturb the vertices shared by C^q and C_b analogously. See in Figure 3.23 an example of how the perturbation of the ith level can look like.

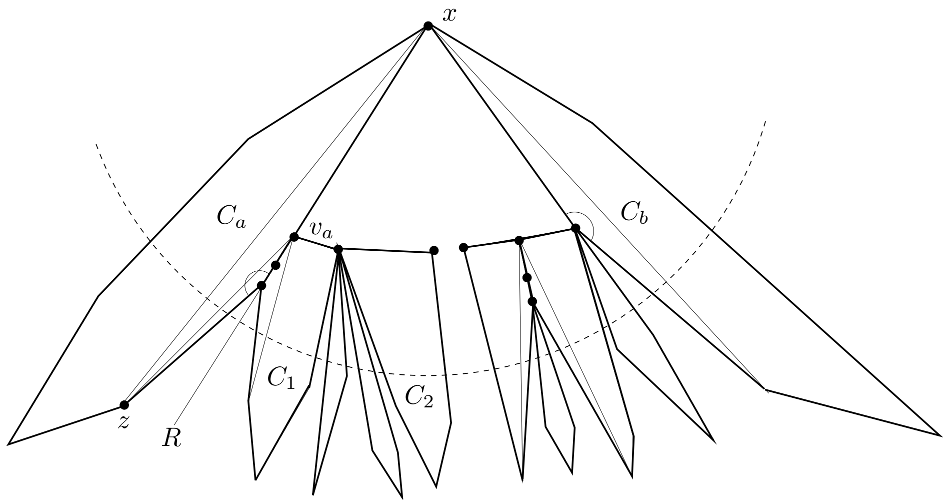

Figure 3.23: Schematic drawing of a perturbation of the ith level. The dashed line represents an arc of $B_{\delta/3}(u)$ and separates the vertices converging to u from the vertices converging to other positions. This arc is actually much closer to x and cigars are much flatter. The marked angles must be strictly convex.

Perturbing the remaining vertices that belong only to one cigar C**.** Finally we place the endpoint vertices of C converging to u within $\mathcal{R}_u$, ensuring that C is strictly convex on those vertices incident to a boundary bar which had non-zero length in $\mathcal{C}_B$. (Note that it can happen that the segment in which we must place some of these vertices is already drawn in its final position.)

Step 3. Obtaining the Perturbation of $\mathcal{C}_{B'}$. Now, the vertices converging to v are all perturbed. We do the same for all the vertices of the convexified drawing of $\mathcal{C}'_B$.

Obtaining the Perturbation of $\mathcal{C}_B$

Any vertex p has been redrawn at most three times: a first time in the convexification of the merged cigars, a second time in the spreading of cigars without vertices in common, and a third time in the final vertex perturbation. Each time it has been perturbed a distance less than $\delta/3$ from its current position. Then, each vertex has been perturbed at most a total distance of δ with respect to its original position in $\mathcal{C}'_B$. Therefore we obtain a δ-perturbation of $\mathcal{C}_B$ where the cigar boundaries are drawn as convex polygons.

Note that the obtained perturbation contains some vertices and bars that were not originally in $\mathcal{C}$.

3.3.3 Second Stage: Perturbing the Interior of the Cigars

Any cigar can be seen as a one-dimensional self-touching configuration with a fixed convex boundary.

If a cigar has chords, the chords divide it into smaller cigars which are also convex, because any chord of a convex polygon divides the polygon into two convex polygons. See Figure 3.24. In this case, we treat independently of each of these smaller cigars, and the chords play the role of boundary bars.

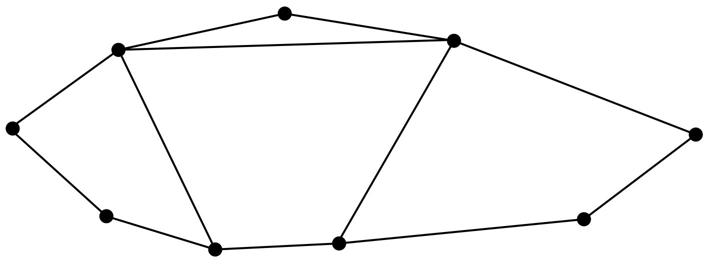

Figure 3.24: Some chords between boundary vertices divide the cigar into smaller convex cigars.

We apply to each cigar independently the method described in Section 3.2 for perturbing one-dimensional configurations. We choose as x-coordinate axis for each cigar the line through its endpoints. Since the cigar boundary vertices are perturbed within δ-disks and the interior vertices remain in the interior of the cigar, by choosing x-coordinates of the interior vertices within δ-disks, the straight-line embedding of the interior of the cigar given by Theorem 3.3 is a δ-perturbation. Note that some vertices of the cigar boundaries may not be strictly convex, but this is not a problem since we have already ensured that they do not have incident chords that could overlap with the other boundary bars.

Thus, we perturb the interior of the cigars independently, obtaining a δ-perturbation of the whole self-touching configuration $\mathcal{C}$. (The vertices and edges artificially added can be removed at the end.) Therefore, Theorem 3.1 is proved.

3.4 Consequences of Theorem 3.1

Corollary 3.1. *Self-touching polygonal chains and cycles are infinitesimally flexible.*

Proof. Let C be a self-touching polygonal chain or cycle. Suppose it is infinitesimally rigid. Then, by Theorem 1.2, it is rigid, and, by Theorem 1.1, it is strongly locked. By Theorem 3.1, there exist a perturbation of C, which is locked by definition of strongly locked self-touching configuration. But this perturbation is a simple polygonal chain or cycle, thus it is unlocked, and we have a contradiction. □

Corollary 3.1 is a first step towards the infinitesimal version of the Carpenter's Rule problem. Although we have proven the existence of an infinitesimal motion, we do not know how to find an initial direction of movement (even for the one-dimensional case) or a global motion, which is still more difficult.