Chapter 18

Union Complexity of Pseudodisks

By Sariel Har-Peled, March 21, 2023[1] Version: 1.2

> I've never touched the hard stuff, only smoked grass a few times with the boys to be polite, and that's all, though ten is the age when the big guys come around teaching you all sorts to things. But happiness doesn't mean much to me, I still think life is better. Happiness is a mean son of a bitch and needs to be put in his place. Him and me aren't on the same team, and I'm cutting him dead. I've never gone in for politics, because somebody always stand to gain by it, but happiness is an even crummier racket, and their ought to be laws to put it out of business.
>
> – – Momo, Emile Ajar.

18.1. Pseudodisks and union complexity

Let $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ be a set of close Jordan curves in the plane. We assume that Γ is a set of ***pseudodisks***. Formally, every two curves of Γ either do not intersect, or intersect in two points. (The case that two curves intersect in a single point in a common tangency point is uninteresting as we can slightly modify the two curves so that they have two intersections). As usual, we assume general position – no three curves intersect in a common point. Observe that being a pseudodisk is a social phenomenal – a curve is one only in the context of a set of other curves. Furthermore, a set of regions being pseudodisk enable them to be arbitrarily "complicated". See Figure 18.1.1.

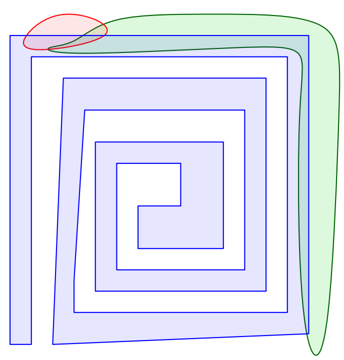

Figure 18.1.1: Pseudodisks.

Let $d_i = d(\gamma_i)$ be the closed region bounded by γ_i, for all i. Let $\mathcal{D} = \{d_1, \ldots, d_n\}$ be the resulting set of closed regions. We are interested in the union $\mathcal{U} = \mathcal{U}(\mathcal{D}) = \cup\mathcal{D} = \cup_{d\in\mathcal{D}} d$. Specifically, we are interested in describing the boundary of $\mathcal{U}$ (i.e., $\partial\mathcal{U}$). This boundary is a collection of close cycles, where each cycle is a sequence of arcs, where each arc is taken from a single curve of Γ. The ***combinatorial***

complexity of the union of the pseudodisks of Γ, denoted by $\langle \Gamma \rangle$, is the total number of edges in these boundary cycles. Our purpose here is to prove that the complexity of the union is $\langle \Gamma \rangle = O(n)$. See Figure 18.1.2.

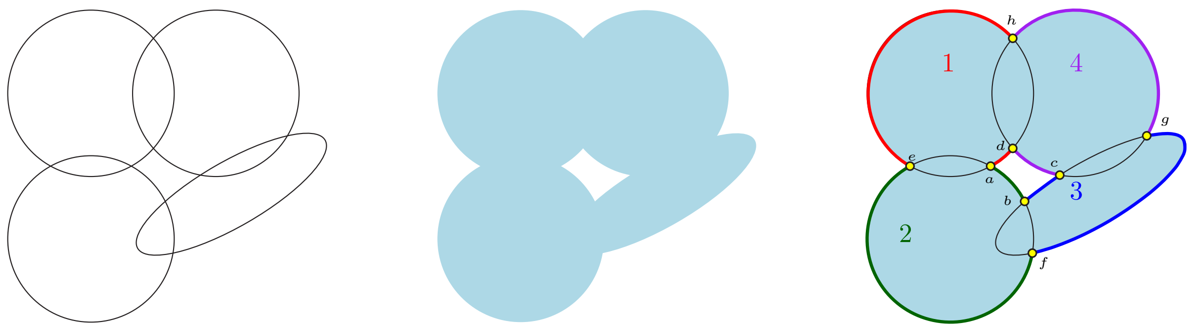

Figure 18.1.2: Here Γ is a set of three circles and an ellipse. The union boundary has two connected components, each one of them is a cycle of edges and vertices. For example, the boundary cycle of the hole can be described as $C = \langle a, 1, d, 4, c, 3, b, 2 \rangle$ (where $1, 2, 3, 4$ are the identifiers of the disks, and a, b, c, d are the vertices) – this is sufficient to recover the edges forming the union. In particular, the combinatorial complexity of C is (say) 4 – the number of vertices appearing in it. In this case, each of the connected components of the boundary of the union form a cycle of combinatorial complexity $O(4)$. Thus $\langle \Gamma \rangle = O(8) = O(1)$.

In the following, we use the shorthand $\Gamma - \gamma = \Gamma \setminus \{\gamma\}$.

Example 18.1.1. Consider a set $\mathcal{R}$ of n disjoint narrow horizontal rectangles, and a set of n narrow vertical rectangles such that every pair of the intersects forming a grid, see Figure 18.1.3. Every pair of rectangles in this example either does not intersect, or intersect in four points. They definitely do not form a collection of pseudodisks. The complexity of the union is $\langle \mathcal{R} \rangle = \Theta(n^2)$.

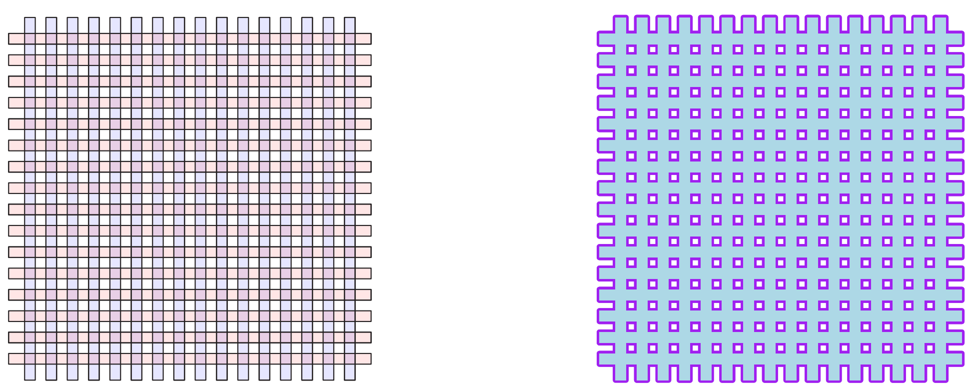

Figure 18.1.3: A grid of n axis-parallel rectangles, and its union. The boundary of the union is made out of $\Theta(n^2)$ polygonal cycles – all internal ones have constant complexity, and the other face has complexity $\Theta(n)$.

18.1.1. A warm-up exercise: The union complexity of disks

As suggested by the above name of pseudodisks, the intuition behind the union complexity being low is disks. So let us first prove this case.

Lemma 18.1.2. *Let $\mathcal{D}$ be a set of n disks in the plane. The union complexity of $\mathcal{D}$ is $O(n)$.*

Proof: Compute the power diagram of $\mathcal{D}$. As a reminder, the power diagram assigns a disk d with center c and radius r, the "distance" function associated with d is $f(p) = \|pc\|^2 - r^2$. Standard lifting argument shows that the power diagram is formed by the lower envelope of planes in three dimensions. As such, the total complexity of the power diagram is $O(n)$, see Figure 18.1.4.

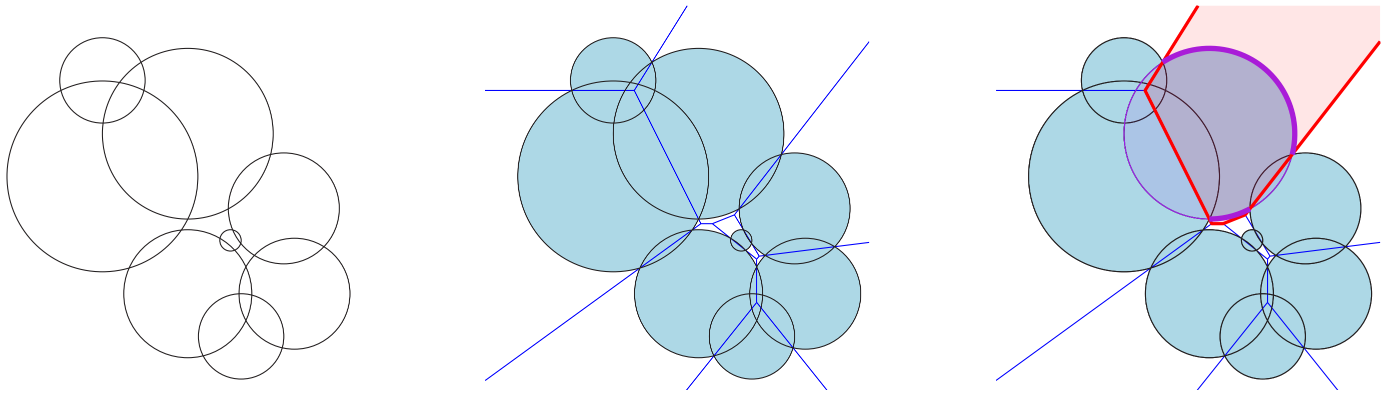

Figure 18.1.4: A set of disks, its power diagram, and single cell in the diagram, and contribution of its disk to the boundary of the union.

The bisector between two disks is a line that passes through the two intersections of their boundary points if they intersect. As such, the cell of a disk $d \in \mathcal{D}$ in the power diagram is a convex polygon Π. Furthermore, for every other disk $d' \in \mathcal{D}$ we have the property that either $d' \cap \Pi = \emptyset$, or $d' \cap \Pi \subseteq d$. Namely, any contribution of d to the union $\cup\mathcal{D}$ is restricted to lie inside Π. The number of times $\partial\Pi$ can intersect ∂d, is bounded by twice the number of edges of Π, Indeed, every edge of Π intersect a circle at most twice.

As such, the contribution of d to the union is bounded by the complexity of its power diagram cell. Since the total complexity of the diagram is $O(n)$, the claim follows. ■

18.2. The proof

The idea of the proof is to redraw the collection of curves so that their union complexity increases while the arrangement they formed becomes simpler. This redrawing process would be done in stages. We start from the obvious – curves that do not contribute to the union are redundant, and can be removed.

Lemma 18.2.1. *A curve $\gamma \in \Gamma$ is* **redundant** *if $\gamma \subseteq \cup_{\sigma\in\Gamma-\gamma} d(\sigma)$ – that is, γ is fully covered by the other pseudodisks of $\mathcal{D}$. We have that $\langle \Gamma - \gamma \rangle \geq \langle \Gamma \rangle$.*

Proof: Observe that no portion of γ appears on $\partial(\cup_{d\in\mathcal{D}} d)$. However, new boundary components might be exposed in $\mathcal{U}(\Gamma)$ as we remove γ from it. That is, $\partial\mathcal{U}(\Gamma) \subseteq \partial\mathcal{U}(\Gamma - \gamma)$. ■

Let $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$, and $\mathcal{D} = \{d(\gamma_i) \mid \gamma_i \in \Gamma\}$. The above lemma implies that one can assume that, in no point in time, we have redundant curves, as one can simply remove them, and continue bounding the complexity of the union on the residual set. Let $\#_3 = \#_3(\Gamma)$ denote the number of triples of regions of $\mathcal{D}$ that intersect, that is

$$\#_3 = \left|\{(i, j, k) \mid i < j < k \text{ and } d_i \cap d_j \cap d_k \neq \emptyset\}\right|.$$

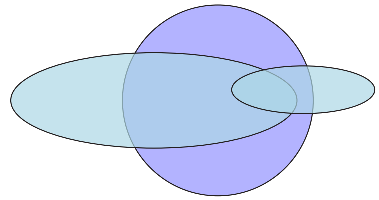

Figure 18.2.1

Lemma 18.2.2. *Let Γ be a set of $n > 2$ pseudodisks. If $\#_3(\Gamma) = 0$, then $\langle \mathcal{U}(\Gamma) \rangle \leq 6n - 12$.*

Proof: For a region $d \in \mathcal{D}(\Gamma)$, observe that $\mathsf{core}(d) = d \setminus \mathcal{U}(\Gamma - \gamma)$ is a non-empty connected set. Indeed, if not, that there must be an intersection of two regions of $\mathcal{D}$ inside d, and that is not possible, as this would imply a triple intersection when including d, see Figure 18.2.1.

The last claim seems intuitively obvious but a proof requires some effort. Indeed, let $d_1, \ldots, d_t$ be all the pseudodisks of $\Gamma - d$ that intersects d. Let $d^i = d \setminus \cup_{\ell=1}^{i} d_\ell$. The claim is that d^i is connected. Indeed, d^{i-1} is connected by induction, and subtracting d_i from it to get d^i can disconnect it only if the boundaries of d_i and d intersect more than twice (impossible!), or d_i intersect some other earlier disk $d_1, \ldots, d_{i-1}$ in the interior of d, which is again preposterous.

Back to the original proof – it is now straightforward to draw a planar graph, where there is an edge between two cores $\iff$ the two corresponding pseudodisks intersect. To this end, for each pseudodisk d place a vertex $v(d)$ in the middle of its core. For two pseudodisks d, d' that intersect, let $L = d \cap d'$ be their common intersection (sometime called a ***lens***), see Figure 18.2.2. Connect $v(d)$ to a point on the boundary of L by a curve that lies in the interior of the core of d. Similarly, connect $v(d')$ to the boundary of L. Finally, connect the two points on the boundary of L via a curve that lies inside L. Clearly, since the lenses are disjoint, the resulting drawing has no intersecting edges, and thus the graph is planar.

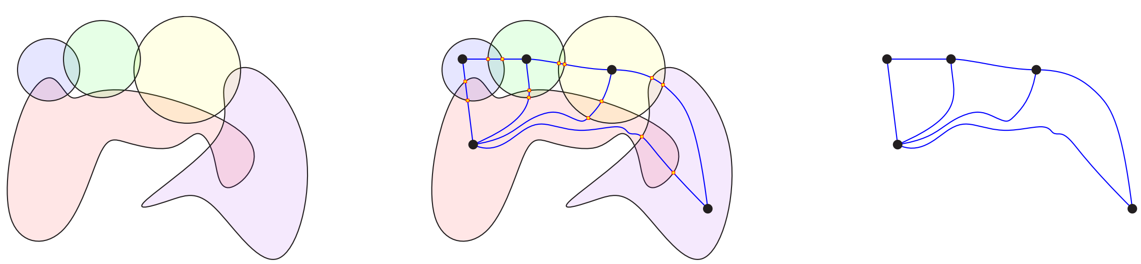

Figure 18.2.2: A set of pseudodisks with no triple intersection, their intersection graph, and the resulting planar graph.

This planar graph has no parallel edges, and every edge can be charged for two vertices on the boundary of the union. It follows by Euler formula that the number of edges is bounded by $3n - 6$, and the number of vertices on the boundary of the union is bounded by $6n - 12$. To see the last step, consider an edge e between d and d'. In the union boundary, this intersection gives rise to two vertices. Since a vertex on the boundary of the union is visited only once (assuming general position), the claim follows. ■

For the next step of the proof we need an intuitive but somewhat heavy hammer (we will not prove it):

Theorem 18.2.3 (Schoenflies theorem). *If $\gamma \subseteq \mathbb{R}^2$ is a simple closed curve, then there is a homeomorphism $f : \mathbb{R}^2 \to \mathbb{R}^2$ such that $f(\gamma)$ is the unit circle in the plane (centered at the origin).*

The above theorem is clear except for what the words used in stating it mean. Being somewhat informal, a ***homeomorphism*** is a continuous (bijective) deformation of the plane. This means that we can deform the plane continuously from its original to its target "drawing" in such away that curves remain curves, and we never introduce new intersections during this process (or remove intersections). Note, that by this continuity requirement, we can *not* do things like flipping the plane (or inversions), since this would require non-intersecting curves to intersect sometime during the process. Two spaces with a homeomorphism between them are ***homeomorphic***, and from a topological viewpoint they are the same.

As such, for our purposes, we can continue the argument on the union complexity of pseudodisks after applying a homeomorphism to the plane and arguing about the complexity of the union of the new pseudodisks.

In the following, let $\mathsf{b}(p, r) = \{q \in \mathbb{R}^2 \mid \|pq\| \le r\}$ be the ***disk*** of radius r centered at p.

Lemma 18.2.4. *Given a collection Γ of n pseudodisks such that $\#_3(\Gamma) > 0$, one can assume there are three curves $\gamma_1, \gamma_2, \gamma_3 \in \Gamma$ with the following properties:*
(I) $d(\gamma_1) \cap d(\gamma_2) \cap d(\gamma_3) \neq \emptyset$.
(II) γ_1 is the unit circle centered at the origin.
(III) For any $i > 1$, if γ_i intersect γ_1, then its intersection with the ring

$$\circledcirc(1, 2) = \mathsf{b}(\mathsf{o}, 2) \setminus \mathsf{b}(\mathsf{o}, 1)$$

is two straight segments that are collinear with the plane, where o denotes the origin.
(IV) For $i > 1$, we have that $\gamma_i \cap \mathsf{b}(\mathsf{o}, 1)$ is either empty or a segment.
Such that the union complexity $\langle\Gamma\rangle$ remains the same.

Proof: (I) follows by renumbering the curves, as $\#_3(\Gamma) > 0$ implies that there are three curves that their interior intersect in a common point.

(II) This is readily implied by Schoenflies theorem.

(III) The intuitive idea is to think about γ_1 has having "thickness" one. Formally, given a drawing of Γ that fulfills the above two properties, and a parameter $\varepsilon \in [0.1]$, consider the mapping $g_\varepsilon : \mathbb{R}^+ \to \mathbb{R}^+$:

$$g_\varepsilon(x) = \begin{cases} x & x \in [0, 1] \\ 1 + (x-1)/\varepsilon & x \in (1, 1+\varepsilon) \\ 2(x - \varepsilon) & x > 1 + \varepsilon. \end{cases}$$

Consider the associated homomorphism $f_\varepsilon : \mathbb{R}^2 \to \mathbb{R}^2$, where $f_\varepsilon(p) = g_\varepsilon\Big(\|p\|\Big)\frac{p}{\|p\|}$. This maps the ring $\circledcirc(1, 1+\varepsilon)$ to the ring $\circledcirc(1, 2)$. In the limit, as $\varepsilon \to 0$, we get thee resulting redrawing of Γ, where any curve that intersect γ_1, intersect $\circledcirc(1, 2)$ in two straight segments collinear with the origin.

(IV) The curve γ_1 is now a circle. If γ_i intersects γ_1, we replace this portion by the segment that connects the two intersection points. Clearly, two such segment intersect on only if the original two curves intersected. Since this redrawing happens inside $d(\gamma_1) = \mathsf{b}(\mathsf{o}, 1)$, it does not effect the union or its complexity. Clearly, the resulting collection of curves is a set of pseudodisks. See Figure 18.2.3. ■

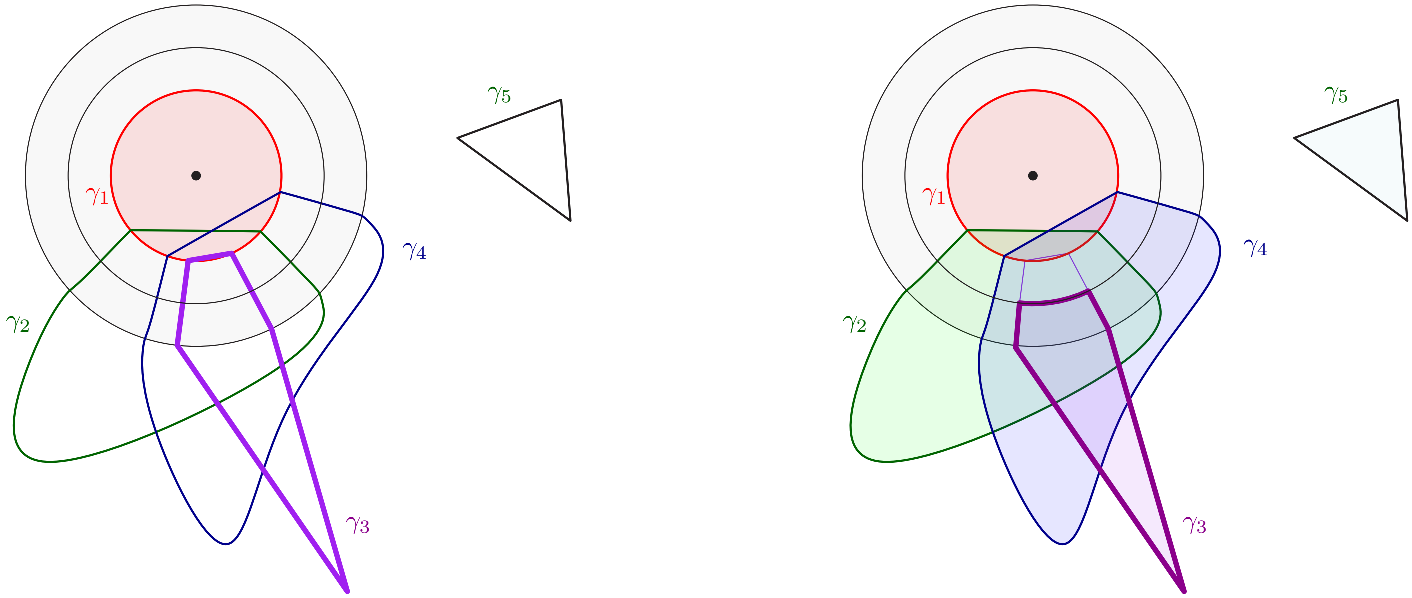

Figure 18.2.3: Left: A set of pseudodisks after we applied Lemma 18.2.4 to it. Right: A curve that is covered by the other pseudodisks on γ_1 can be "snapped" out, in the process reducing the number $\#_3$ of triple intersections.

Lemma 18.2.5. *Given a set Γ of pseudodisks Γ with $\#_3(\Gamma) > 0$, one can replace it with a set of pseudodisks Γ', such that $\#_3(\Gamma') < \#_3(\Gamma)$ and the $\langle \Gamma' \rangle \geq \langle \Gamma \rangle$.*

Proof: We apply the redrawing of Γ as specified by Lemma 18.2.4 so that γ_1 is the union circle $\gamma_1 = \partial \mathsf{b}(\mathsf{o}, 1)$, and there at least two other curves γ_2 and γ_3 that intersect γ_1 (there might be many more). let $d_i = d(\gamma_i)$, for all i. Importantly, $d_1 \cap d_2 \cap d_3 \neq \emptyset$. We assume that properties (I)-(IV) of Lemma 18.2.4 hold for this drawing.

For every curve γ_i that intersect γ_1, its region d_i intersects γ_1 in a circular arc, and let $I_i = \gamma_1 \cap d_i$ denote this circular arc. If I_3 is covered by the union of pseudodisks (ignoring d_1), then we redraw γ_3, so that its interior inside $\mathsf{b}(\mathsf{o}, 2)$ is replaced by the (scaled) arc $(3/2)I_3$ together with two segments that connect it to the outside of $\mathsf{b}(\mathsf{o}, 2)$. It is easy to verify that the new set of curves is still a set of pseudo disks – see Figure 18.2.3 and Figure 18.2.4. Furthermore, the union of pseudodisks have not changed, but γ_1 and γ_3 no longer intersect (bye bye, and thanks for all the intersections!). Namely, for the resulting set of pseudodisks Γ', we have that $\#_3(\Gamma') < \#_3(\Gamma)$ and $\langle \Gamma' \rangle = \langle \Gamma \rangle$.

We can apply the above argument if any other arc I_j is covered by the union of the remaining arcs. So assume this does not happen – every arc I_i has at least one point on γ_1 that no other curve covers. In this case, we do the same thing – we redraw γ_3 by replacing the portion between the endpoints of $(3/2)I_3$ on γ_3 by $(3/2)I_3$. It is again easy to verify that the resulting set of curves Γ' is a set of pseudo disks. However, interestingly, the union $\mathcal{U}(\Gamma')$ now has a hole that was previously covered by d_3, see Figure 18.2.4. Namely, the union complexity of the new set of pseudodisks is strictly larger. We conclude that $\#_3(\Gamma') < \#_3(\Gamma)$ and $\langle \Gamma' \rangle > \langle \Gamma \rangle$. ∎

Theorem 18.2.6. *Given a set Γ of n pseudodisks in the plane, their union complexity is at most $6n - 12$.*

Proof: As long as $\#_3(\Gamma) > 0$, we apply Lemma 18.2.5 to get a new family of curves of the same size that has the same or bigger union complexity, but fewer intersections of triples of regions. After a finite number of such redraws we end up with a family of pseudodisks with no triple intersection. But then the bound follows from Lemma 18.2.2. ∎

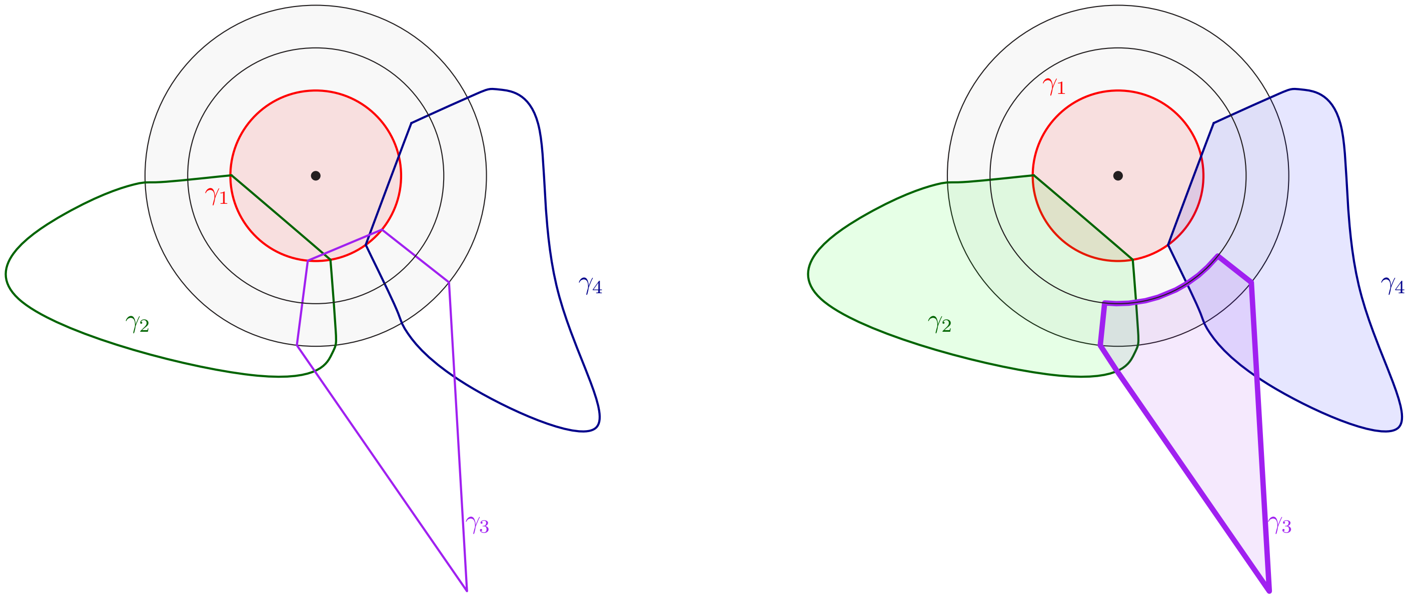

Figure 18.2.4: Reducing the number of triple intersecting by snapping one of the pseudodisks outside.

18.3. Bibliographical notes

Our presentation more or less follows the original proof of Kedem *et al.* [KLPS86]. Arguably our presentation is somewhat simpler.

References

[KLPS86] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.*, 1: 59–71, 1986.