

(δ, ε) -ball approximation of a shape: definition and complexity

Dominique Attali · Tuong-Bach
Nguyen · Isabelle Sivignon

the date of receipt and acceptance should be inserted later

Abstract Given a set S in \mathbb{R}^n , a (δ, ε) -ball approximation of S is defined as a collection of balls that covers the morphological erosion of S (by a ball of radius ε) and remains inside the morphological dilation of S (by a ball of radius δ). We study the problem of computing a (δ, ε) -ball approximation when S is itself defined as a finite union of balls. This problem relates to geometric set cover problems but is however different in nature. It offers a new framework for reducing the size of a collection of balls while controlling both the inner and outer distance to the shape. We prove that computing a (δ, ε) -ball approximation of minimum cardinality is NP-complete for $n = 2$. Along the way, we study the boundary of unions of disks and their erosion, for which we derive a computational description.

Keywords shape approximation · union of balls · medial axis · covering · morphological erosion and dilation · complexity

Mathematics Subject Classification (2000) 68U05 · 52C99

Research partially supported by the French “Agence nationale pour la Recherche” under grants ANR-13-BS01-0008 TopData. Also partially supported by the French program “Grenoble Innovation Recherche” 2014 from Université Grenoble Alpes and Grenoble INP.

Dominique Attali · Tuong-Bach Nguyen · Isabelle Sivignon
Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-Lab, 38000 Grenoble, France
E-mail: dominique.attali, tuong-bach.nguyen, isabelle.sivignon@gipsa-lab.grenoble-inp.fr

1 Introduction

Context. In many applications, shapes are often represented as the union of a finite set of balls. Such a representation is used in biochemistry to model molecules [17, 10] or in computer graphics to detect collisions between objects [5]. Unions of balls have become ubiquitous representations of shapes, largely due to the existence of provably good conversion algorithms that allow to build them from other representations such as point clouds [2], polygonal meshes [15, 24, 27] or digital shapes [14]. Conversion algorithms generally output very large (and sometimes infinite) collections of balls. This may be the case when conversion algorithms aim at representing the shape as the union of its medial balls. A ball b is said to be *medial* for a given shape if every ball that contains b and is contained in the shape is equal to b . Assuming that balls are closed and the shape is compact, the union of medial balls provides an exact description of the shape, whose size however may be infinite. The centers of medial balls form the *medial axis*, whose stability and computation (either exact or approximated) have been widely studied; see [3] for a state-of-the-art report. For computationally demanding tasks such as the simulation of physical processes [19], shape interpolation or shape matching [30, 8], it is desirable that the collection of balls that describes a shape has a small size while its union still provides an accurate approximation of the shape. The quality of the approximation is usually measured by various quantities such as the Hausdorff distance or the difference in volume between the shape and its representation. In other situations such as collision detection, it may be important that the union of balls satisfies some *geometric constraints*. For instance, we may want the balls to cover the shape or at least a prescribed set of points.

Ball approximations. In this work, we introduce a novel way of constraining a collection of balls so that its union remains close to the shape. The idea is to force balls to cover a subset of the shape while remaining contained in a superset of the shape, hence allowing balls neither to cover exactly nor to be contained perfectly in the shape. To explain this more precisely, we need some notation and definitions. Unless stated otherwise, all balls considered in the paper are closed. We denote by $b(x, r)$ the (closed) ball centered at x with radius r . Given a subset $S \subseteq \mathbb{R}^n$ and a real number $r \geq 0$, we recall that the *dilation* of S by a ball of radius r is $S^{\oplus r} = \cup_{x \in S} b(x, r)$ and the *erosion* of S by a ball of radius r is $S^{\ominus r} = \{x \mid b(x, r) \subseteq S\}$. Note that if S is closed, then both $S^{\ominus r}$ and $S^{\oplus r}$ are closed. We shall refer to $S^{\ominus r}$ and $S^{\oplus r}$ as the r -erosion and r -dilation.

Definition 1 Let $S \subseteq \mathbb{R}^n$, $\varepsilon \geq 0$ and $\delta \geq 0$. A collection of balls \mathcal{B} is a (δ, ε) -ball approximation of S if $S^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B} \subseteq S^{\oplus \delta}$.

Let us make a few remarks. Consider a (δ, ε) -ball approximation \mathcal{B} of S ; see Figure 1 for an example. Setting $B = \bigcup \mathcal{B}$, we thus have $B \subseteq S^{\oplus \delta}$ and $(S^{\ominus \varepsilon})^{\oplus \varepsilon} \subseteq B^{\oplus \varepsilon}$. It follows that if $S = (S^{\ominus \varepsilon})^{\oplus \varepsilon}$, then $B \subseteq S^{\oplus \delta}$ and $S \subseteq B^{\oplus \varepsilon}$, or equivalently the Hausdorff distance between B and S is smaller than or




Fig. 1: A shape (in blue), its ε -erosion (in orange), its δ -dilation (bounded by the dashed blue curve) and a (δ, ε) -ball approximation (the three green balls).

equal to the maximum of δ and ε , *i.e.* $d_H(S, B) \leq \max\{\delta, \varepsilon\}$. The condition $S = (S^{\ominus\varepsilon})^{\oplus\varepsilon}$ can be seen as some regularity¹ assumption on S . Indeed, the equality $S = (S^{\ominus\varepsilon})^{\oplus\varepsilon}$ is satisfied if every point of S can be covered by a ball of radius ε which is contained in S . In particular, the condition holds whenever all medial balls have a radius larger than or equal to ε , which in turn holds whenever the reach of S is larger than or equal to ε . We recall that the reach of S was first introduced by Federer [18] and is the infimum of distances between points in the medial axis of S and points in the complement of S . Denoting the reach of S by $\text{reach}(S)$, we summarize our findings in the following remark:

Remark 1 Let \mathcal{B} be a (δ, ε) -ball approximation of S . If $\text{reach}(S) \geq \max(\delta, \varepsilon)$ then $d_H(\bigcup \mathcal{B}, S) \leq \max(\delta, \varepsilon)$.

Remark 2 S possesses (δ, ε) -ball approximations of finite cardinality whenever S is compact and $\delta + \varepsilon > 0$.

Suppose S is compact and $\delta + \varepsilon > 0$. To build a finite (δ, ε) -ball approximation of S , consider the collection of open balls with radius $\delta + \varepsilon$ centered at points in $S^{\ominus\varepsilon}$. By construction, the collection covers $S^{\ominus\varepsilon}$ which is compact because S is compact, implying that the collection has a finite subcover. By taking the closure of balls in that subcover, we get a collection of closed balls that are contained in $S^{\oplus\delta}$ and thus form a finite (δ, ε) -ball approximation of S .

¹ A set S is called *r-regular* if $S = (S^{\ominus r})^{\oplus r} = (S^{\oplus r})^{\ominus r}$ [32]. In other words, the shape S is *r-regular* if it does not change neither under a morphological opening nor a morphological closing by a ball of radius r .

Main result. In this paper, we are interested in a discrete variant of the following problem:

Problem 1 (Ball approximation) Given a subset $S \subseteq \mathbb{R}^n$, two non-negative real numbers δ and ε , find a (δ, ε) -ball approximation of S of minimum cardinality.

Our main result is that a discrete decision version of this problem is NP-complete. To state the problem, we call any ball whose center has rational coordinates and whose radius is rational *rational*. We call any collection of rational balls *rational*.

Problem 2 (Rational ball approximation) Given a finite collection of rational balls \mathcal{S} in \mathbb{R}^n , two non-negative rational numbers δ and ε and an integer $k > 0$, answer the following question: does $\bigcup \mathcal{S}$ has a rational (δ, ε) -ball approximation with cardinality lower than or equal to k ?

The main result of this paper is that this problem is already difficult in dimension 2:

Theorem 1 *The rational ball approximation problem is NP-complete in \mathbb{R}^2 .*

Related works. As mentioned above, any shape S can be exactly described by the union of its medial balls \mathcal{M} . A (δ, ε) -ball approximation of S of minimum cardinality is a collection of balls \mathcal{B} such that $|\mathcal{B}| \leq |\mathcal{M}|$. Indeed, note that, by Definition 1, \mathcal{M} itself is a (δ, ε) -ball approximation of S . Therefore, solving Problem 1 offers a new framework for *simplifying unions of balls* with a control of the distance to the shape. Various algorithms have been proposed to compute simpler collections of balls. A family of algorithms was designed to deal with the sensitivity of the medial axis to small perturbations on the shape S : the general idea is to compute a geometric quantity for each medial ball and prune the balls whose quantity is below a given threshold (see for instance [12, 3, 11, 16] to cite only a few). A slightly different approach has been presented in [27] with the notion of scale axis, where balls of the medial axis are first inflated, the medial axis updated (resulting in a simplification) and its balls deflated. Last, hierarchical structures have also been proposed to obtain simpler collections of balls. In the context of collision detection, sphere-trees are built by iteratively merging and/or displacing balls while ensuring that the simpler set of balls includes the boundary of the input one [23, 5, 7].

Note that in all the approaches mentioned above, the simpler collection of balls results in a new shape that is either a subset or a superset of the input shape. In that respect, the (δ, ε) -ball approximation problem offers a more general setting that encompasses these approaches: indeed, when $\delta = 0$, a (δ, ε) -ball approximation results in a new shape that is a subset of the input shape; conversely, when $\varepsilon = 0$, the new shape is a superset of the input shape.

The (δ, ε) -ball approximation problem is also closely related to *geometric set cover* problems. Following [1] or [21], let $\Sigma = (K, R)$ be a finite geometric range space where K is a *finite* set of points in \mathbb{R}^d and R , the range, is a

finite family of simple shaped regions, like for instance rectangles, or disks. A subset $C \subseteq R$ is a cover of K if any point of K belongs to at least one element of C . Computing a set cover of minimum cardinality is known as the set cover problem. The corresponding decision problem is well-known to be NP-complete, and various approximation algorithms have been proposed, from the simple greedy one [13, 25], to more convoluted algorithms using ε -nets or multiplicative weights methods [6, 21, 1]. Several variations of the geometric set cover problem have also been introduced and investigated, notably when the range R is a family of balls. Instead of minimizing the number of balls used to cover the whole set K , a related problem is to cover the set K at best using a prescribed number of balls (see for instance a greedy algorithm in [10]). For more applicative purposes, especially in sensor networks, the balls of the range R can be of prescribed centers, but undefined radii, and the goal is to find a cover that minimizes a function of the radii of the balls [31].

The key difference between (δ, ε) -ball approximation and set cover problems lies in the assumption of a finite range space for the latter one. If the shape K to cover and the range R are infinite, another family of results exists for which the range is the (infinite) set of congruent balls of a given radius. Results on that matter go back as far as [26], that establishes a relation between the (Lebesgue) measure of a bounded set of points K in \mathbb{R}^2 and the number of congruent disks of radius r necessary to cover K . When K is a simple geometric shape (disk, square, equilateral triangle), results on the maximum size (radius, side length) of a shape K that can be covered by n unit disks are documented in [33].

Paper overview. We establish Theorem 1 by successively proving that:

- The rational ball approximation problem in \mathbb{R}^2 is in NP. Calling a closed ball in \mathbb{R}^2 a disk, this means that one can verify in polynomial time that a given rational collection of disks is indeed a solution of the problem. For this, we start by studying in Section 2 the boundary of the erosion of a union of disks, for which we give a full computational characterization. This result is then used in Section 3 to design a polynomial time algorithm to check that a given collection of disks covers the ε -erosion $S^{\ominus\varepsilon}$ and is included in the δ -dilation $S^{\oplus\delta}$ using arrangements of circles.
- The rational ball approximation problem in \mathbb{R}^2 is NP-hard. We prove this through a reduction from the well-known NP-complete vertex cover problem in Section 4.

Preliminaries. Given a subset $X \subseteq \mathbb{R}^n$, we let \overline{X} denote the closure of X , $\overset{\circ}{X}$ the interior of X and ∂X the boundary of X .

2 Studying boundaries

Let \mathcal{S} be a finite collection of disks in \mathbb{R}^2 and ε a non-negative real number. The purpose of this section is to provide a computational description of the

boundary of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$. For this, we introduce in Section 2.2 a superset of the boundary called the *wavefront* and characterize which part of the wavefront belongs to the boundary. In Section 2.3, we decompose the wavefront into elements that are easy to compute, based on the decomposition of $\partial(\bigcup \mathcal{S})$ into vertices and edges described in Section 2.1.

2.1 Boundary of a finite union of disks

Given a finite set of (closed) disks \mathcal{S} in \mathbb{R}^2 , we decompose the boundary of the union $\bigcup \mathcal{S}$ into a finite set of elements that we call *vertices* and *edges*. For this, we associate to each point $u \in \mathbb{R}^2$ the subsets of disks that have the property to represent $\bigcup \mathcal{S}$ locally around u (see Definition 2) and rely on those subsets to distinguish between vertices or edge points of $\partial(\bigcup \mathcal{S})$. All the section is in \mathbb{R}^2 , except Definition 2 and Property 1 which are stated in \mathbb{R}^n .

Definition 2 Given a subset $X \subseteq \mathbb{R}^n$ and a point $u \in \mathbb{R}^n$, we say that a set of balls \mathcal{B} in \mathbb{R}^n *represents X locally around u* if:

- (i) for all $b \in \mathcal{B}$, $u \in b$;
- (ii) there exists $r > 0$ such that $X \cap b(u, r) = \bigcup \mathcal{B} \cap b(u, r)$.

We first establish that such a set \mathcal{B} exists and is finite when X is the union of a finite collection of balls:

Property 1 Let \mathcal{S} be a collection of balls in \mathbb{R}^n . The set $\mathcal{S}_u = \{b \in \mathcal{S} \mid u \in b\}$ represents $\bigcup \mathcal{S}$ locally around u , for all $u \in \mathbb{R}^n$.

Proof For any ball $b \in \mathcal{S}$, either $u \in b$ or the distance between u and b is positive. Because \mathcal{S} is finite, $\mathcal{S} \setminus \mathcal{S}_u$ is also finite and we let r be the smallest distance between u and a ball in $\mathcal{S} \setminus \mathcal{S}_u$. We have $r > 0$, and by construction $b(u, r) \cap \bigcup \mathcal{S} \subseteq \bigcup \mathcal{S} \setminus \bigcup (\mathcal{S} \setminus \mathcal{S}_u) \subseteq \bigcup \mathcal{S}_u \subseteq \bigcup \mathcal{S}$. Thus $\bigcup \mathcal{S} \cap b(u, r) = \bigcup \mathcal{S}_u \cap b(u, r)$. \square

Let $S = \bigcup \mathcal{S}$ be the union of a finite collection of disks \mathcal{S} in \mathbb{R}^2 . We are now ready to define vertices and edges of ∂S . We say that a point $u \in \partial S$ is a *vertex* if all sets \mathcal{B} that represent S locally around u have cardinality greater than 2, that is, $|\mathcal{B}| \geq 2$. Letting \mathcal{V} be the collection of vertices of ∂S , we call any maximal connected component of $\partial S \setminus \mathcal{V}$ an *edge* of ∂S . Then, for any point u on an edge, there exists a disk $b \in \mathcal{S}$ such that $\{b\}$ represents S locally around u . Notice that $\{b\}$ is the unique set of cardinality 1 that represents S locally around u , because if two disks b_1 and b_2 share a common boundary point u and coincide in a neighborhood of u , then they are equal. We now show that this disk b is the same for all boundary points in a small neighborhood of u :

Property 2 Consider $u \in \partial S \setminus \mathcal{V}$ and let $\{b\}$ that represents S locally around u . Then, there exists $r > 0$ such that $\{b\}$ is the unique cardinal-wise minimum element that represents S locally around v for all $v \in \mathring{b}(u, r) \cap \partial S$.

Proof Let $r > 0$ such that $S \cap b(u, r) = b \cap b(u, r)$. Let $\rho_v = r - \|u - v\|$ and observe that $b(v, \rho_v) \subseteq b(u, r)$. For all $v \in \mathring{b}(u, r) \cap S$, we thus have that $S \cap b(v, \rho_v) = b \cap b(v, \rho_v)$ and it follows that for all points $v \in \mathring{b}(u, r) \cap S$, $\{b\}$ represents S locally around v . In particular, if $v \in \mathring{b}(u, r) \cap \partial S$, then $v \in \partial S \setminus \mathcal{V}$, and uniqueness and minimality stem from the definition of an edge. \square

We deduce immediately that any edge is supported by the boundary of a unique disk in \mathcal{S} . Formally:

Property 3 Let e be an edge of ∂S . Then there is a unique $b \in \mathcal{S}$ such that (1) $e \subseteq \partial b$; (2) $\{b\}$ represents S locally around each point of e .

Proof Consider $f : \partial S \setminus \mathcal{V} \rightarrow \mathcal{S}$ that associates to each point $x \in \partial S \setminus \mathcal{V}$ the unique inclusion-minimum set that represents S locally around x . By Property 2, f is locally constant and therefore constant over each connected component of $\partial S \setminus \mathcal{V}$. \square

2.2 Boundary of the erosion and wavefront

Recall that \mathcal{S} is a finite collection of disks in \mathbb{R}^2 and let $S = \bigcup \mathcal{S}$ be the union of those disks. In this section, we relate the boundary of the erosion $\partial S^{\ominus \varepsilon}$ to a (possibly self-intersecting) curve that we call the *wavefront* of ∂S at time ε . Roughly speaking, the wavefront at time ε is obtained by moving with speed 1 some of the points on the boundary of S a given distance ε down the inward-pointing normals. Formally, the contribution of a point $u \in \partial S$ to the wavefront at time ε is:

$$w_\varepsilon(u) = \{x \in \partial b(u, \varepsilon) \mid \exists r > 0, b(x, \varepsilon) \cap b(u, r) \subseteq S\}.$$

The wavefront of ∂S is obtained by taking the union of all contributions $w_\varepsilon(\partial S) = \bigcup_{u \in \partial S} w_\varepsilon(u)$ (see Figure 2.b). It is not difficult to see that the wavefront contains the boundary of the erosion (see Figure 2.a). Indeed, if a point x belongs to the boundary of the erosion $\partial S^{\ominus \varepsilon}$, then it is the center of a disk $b(x, \varepsilon) \subseteq S$ whose boundary meets the boundary of S at some point u . It follows that $x \in \partial b(u, \varepsilon)$ and thus belongs to the contribution $w_\varepsilon(u)$ of point u . The goal of this section is to establish that the boundary of the erosion $\partial S^{\ominus \varepsilon}$ is the set of points on the wavefront $w_\varepsilon(\partial S)$ at distance ε or more to vertices of ∂S (see Figure 2.c):

Theorem 2 *Let \mathcal{V} be the vertex set of ∂S . For $\varepsilon \geq 0$, we have*

$$\partial S^{\ominus \varepsilon} = w_\varepsilon(\partial S) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ.$$

Proof The inclusion $\partial S^{\ominus \varepsilon} \subseteq w_\varepsilon(\partial S) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ$ is obvious. Let us prove the converse inclusion $w_\varepsilon(\partial S) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ \subseteq \partial S^{\ominus \varepsilon}$. Take a point x in $w_\varepsilon(\partial S) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ$. This implies that there is a point $u \in \partial S \cap \partial b(x, \varepsilon)$ and $r > 0$ such that $b(x, \varepsilon) \cap b(u, r) \subseteq S$. Consider the half-line $[ux)$ with origin u and passing




Fig. 2: (a) A shape defined as the union of three (blue) disks and its erosion (in orange); (b) wavefront of the boundary (red curve); (c) arrangement of circles used to decompose the erosion into cells (see Section 3.1). Three of the circles are centered on the (red) vertices of the boundary. The vertices of the arrangement are the black dots.

through x and for $t \geq 0$, let $z_t \in [ux)$ such that $d(u, z_t) = t$. Denote by b_t the disk centered at z_t and whose bounding circle goes through u , see Figure 3.a. By construction, the family of disks $\{b_t\}$ is nested and their boundaries have a common intersection point u . Let us prove that $b_\varepsilon \subseteq S$, which entails $x = z_\varepsilon \in \partial S^{\ominus\varepsilon}$.

First, observe that for $t > 0$ small enough, we have $b_t \subseteq b(x, \varepsilon) \cap b(u, r) \subseteq S$ and $\partial b_t \cap \partial S = \{u\}$. Hence, as we increase continuously t from 0 to $+\infty$, the center z_t moves on the half-line $[ux)$ and the ball b_t which is contained in S for small values of t goes eventually outside S because S is bounded and b_∞ is not. Let T the largest t for which $b_t \subseteq S$. For all $t < T$, we have $\partial b_t \cap \partial S = \{u\}$. We have $T > 0$ and suppose for a contradiction that $T < \varepsilon$. We claim that b_T meets the boundary of S in another point $y \neq u$ as in Figure 3.b. Indeed, for $t > T$, let $y_t \in b_t \cap S^c$. The sequence (y_t) admits a limit point $y \in \partial b_T$. If we had $y = u$, then for t close enough to T , we would have $b_t \subseteq b_\varepsilon$ and $y_t \in b_\varepsilon \cap b(u, r) \subseteq S$, which is impossible. Thus $y \neq u$. Since by hypothesis $b_\varepsilon = b(x, \varepsilon)$ does not contain any vertex of ∂S in its interior, neither does $b_T \subseteq b_\varepsilon$ and y belongs to some edge e of ∂S . This edge is supported by the boundary of a ball $b \in \mathcal{S}$ that represents S locally around y . It follows that b and b_T are tangent at y with $b_T \subseteq b$. Since $b \subseteq S$ cannot contain $u \in \partial S$ in its interior, we get $b = b_T \subseteq b_\varepsilon$ and the edge e is either the punctured circle $\partial b_T \setminus \{u\}$ or the whole circle ∂b_T depending on whether u is a vertex of ∂S or not. In both cases, $\partial b_T \subseteq \partial S$. Since b_T is contained in b_ε which is contained in S locally around u , the only possibility is that $b_T = b_\varepsilon$, yielding $T = \varepsilon$ and therefore a contradiction. \square




Fig. 3: Notation for the proof of Theorem 2. (a) Construction of z_t and b_t ; (b) second boundary point y .

2.3 Computational description of the wavefront

The goal of this section is to give a computational description of the wavefront. In particular, we shall see that the wavefront is composed of circular arcs; convex circular arcs are contributed by edges of ∂S while concave circular arcs are contributed by vertices of ∂S . To describe those arcs, let us start by studying the contribution of a point $u \in \partial S$ to the wavefront at time ε . We show that this contribution can be expressed as the intersection of $\partial b(u, \varepsilon)$ and an affine conical hull with apex u . Consider a finite collection of points $C = \{c_1, c_2, \dots, c_k\}$ in \mathbb{R}^2 . We recall that the affine conical hull with support C and apex $u \in \mathbb{R}^2$ is

$$\text{coni}(u, C) = \left\{ u + \sum_{i=1}^k \lambda_i (c_i - u) \mid \lambda_i \in \mathbb{R}, \lambda_i \geq 0 \right\}.$$

For technical reasons, we shall need to exclude some of the points on the boundary of affine conical hulls. Writing $[ab]$ for the closed segment with endpoints a and b , this motivates the following definition:

$$\text{coni}^*(u, C) = \text{coni}(u, C) \cup \left(\bigcup_{c \in C} [uc] \right).$$

Theorem 3 *Let $S = \bigcup \mathcal{S}$ be a finite union of disks and consider $u \in \partial S$. Let $\mathcal{B} \subseteq \mathcal{S}$ that represents S locally around u , and C the centers of disks in \mathcal{B} . We have*

$$w_\varepsilon(u) = \text{coni}^*(u, C) \cap \partial b(u, \varepsilon).$$

An illustration is available in Figure 4. Before proving the theorem, we apply it to describe for each element on the boundary of the union $\bigcup \mathcal{S}$ its contribution to the wavefront:




Fig. 4: A shape represented locally around u by two (blue) disks b_1 and b_2 . The contribution of u to the wavefront is the (black) circular arc, closed at x_1 and open at x_2 . By Theorem 3, it is the intersection of the (orange) affine conical hull and $\partial b(u, \varepsilon)$. The endpoint x_1 belongs to the wavefront (unlike the endpoint x_2) because the disk $b(x_1, \varepsilon)$ is locally included in the shape around u (unlike the disk $b(x_2, \varepsilon)$).

Contribution of vertices. Assume that v is a vertex of ∂S . Then any set \mathcal{B} that represents S locally around v is such that $|\mathcal{B}| \geq 2$, and $w_\varepsilon(v)$ is a concave circular arc in the general case. In the special case when \mathcal{S} has two disks b_1 and b_2 meeting tangentially at v and whose pair $\mathcal{B} = \{b_1, b_2\}$ represents locally S around v , then $w_\varepsilon(v) = [c_1 c_2] \cap \partial b(v, \varepsilon)$, where c_i is the center of b_i for $i \in \{1, 2\}$. It follows that $w_\varepsilon(v)$ is composed of zero, one or two degenerate circular arcs (each reduced to a point), depending on whether zero, one or two disks among b_1 and b_2 have radii larger than or equal to ε .

Contribution of edges. Let e be an edge of ∂S . By Property 3, there is a disk $b \in \mathcal{S}$ whose boundary supports e and which represents S locally around each point $u \in e$. Letting c be the center of b and plugging in Theorem 3 the equality $C = \{c\}$ yields $w_\varepsilon(u) = [uc] \cap \partial b(u, \varepsilon) = [uc] \cap \partial b^{\ominus \varepsilon}$ for all $u \in e$. It follows that the contribution $w_\varepsilon(e) = \bigcup_{u \in e} w_\varepsilon(u)$ of an edge e is either empty if b has a radius smaller than ε or is the convex circular arc obtained by projecting e orthogonally onto $b^{\ominus \varepsilon}$ otherwise.

The remainder of this section is dedicated to the proof of Theorem 3. The proof relies on two technical lemma that we now state. Consider a family of disks that all share a common boundary point u . The first technical lemma says that if the center of one of the disk b_0 is contained in the open affine conical hull formed by the centers of other disks, then b_0 is contained in the union of the other disks locally around u . The second technical lemma gives a slightly weaker version of the converse.

Lemma 1 *Let $u \in \mathbb{R}^2$ and consider k points $c_1, c_2, \dots, c_k \in \mathbb{R}^2$. Let $c_0 \in \text{coni}(u, \{c_i, 1 \leq i \leq k\})$. For $0 \leq i \leq k$, let b_i be the disk centered at c_i whose bounding circle goes through u . Then, there exists $r > 0$ such that $b_0 \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$.*




Fig. 5: Notation for the proof of Lemma 1. The bold curve represents the set of points at distance $R_{\mathbf{e}}$ from the origin in directions \mathbf{e} such that $\langle \mathbf{e}, c_0 - u \rangle \geq 0$.

Proof Consider the closed half-space H_0 through u containing b_0 . We prove the theorem by establishing that $H_0 \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$, for some $r > 0$.

Notice that for all unit vectors \mathbf{e} and real numbers ρ , the point $u + \rho\mathbf{e}$ belongs to b_i if and only if $0 \leq \rho \leq 2\langle \mathbf{e}, c_i - u \rangle$. It follows that $u + \rho\mathbf{e}$ belongs to the union $\bigcup_{i=1}^k b_i$ if and only if $0 \leq \rho \leq R_{\mathbf{e}}$ where $R_{\mathbf{e}} = 2 \max_{1 \leq i \leq k} \langle \mathbf{e}, c_i - u \rangle$; see Figure 5.

Let $r = \inf\{R_{\mathbf{e}} \mid \langle \mathbf{e}, c_0 - u \rangle \geq 0, \|\mathbf{e}\| = 1\}$. We next prove that $r > 0$. Consider a unit vector \mathbf{e} with $\langle \mathbf{e}, c_0 - u \rangle \geq 0$. Because c_0 lies in the interior of $\text{coni}(u, \{c_i, 1 \leq i \leq k\})$, there exist k real positive numbers $\lambda_1, \lambda_2, \dots, \lambda_k$ such that $c_0 = u + \sum_{i=1}^k \lambda_i(c_i - u)$; see Appendix B. By linearity,

$$0 \leq \sum_{i=1}^k \lambda_i \langle \mathbf{e}, c_i - u \rangle \leq \left(\sum_{i=1}^k \lambda_i \right) \max_{1 \leq i \leq k} \langle \mathbf{e}, c_i - u \rangle.$$

Assume for a contradiction that $\max_{1 \leq i \leq k} \langle \mathbf{e}, c_i - u \rangle \leq 0$. Since $\lambda_i > 0$ for all $1 \leq i \leq k$, this would imply that $\langle \mathbf{e}, c_i - u \rangle = 0$ for all $1 \leq i \leq k$, showing that all vectors $c_i - u$ are collinear. Hence, the affine conical hull $\text{coni}(u, \{c_i, 1 \leq i \leq k\})$ would be either a line or a half-line, contradicting our hypothesis that its interior $\overset{\circ}{\text{coni}}(u, \{c_i, 1 \leq i \leq k\})$ contains c_0 and is non-empty.

To conclude, consider a point $x \in H_0 \cap b(u, r)$. Let $\rho = \|x - u\|$ and \mathbf{e} the unit vector such that $x = u + \rho\mathbf{e}$. Because $x \in H_0$, $\langle \mathbf{e}, c_0 - u \rangle \geq 0$. Because $x \in b(u, r)$, $0 \leq \rho \leq r \leq R_{\mathbf{e}}$. Hence, $x \in \bigcup_{i=1}^k b_i$ and $H_0 \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$. \square

Lemma 2 *Let $u \in \mathbb{R}^2$ and $k+1$ points $c_0, c_1, \dots, c_k \in \mathbb{R}^2 \setminus \{u\}$. For $0 \leq i \leq k$, let b_i be the disk centered at c_i whose bounding circle goes through u . If there exists $r > 0$ such that $b_0 \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$, then $c_0 \in \text{coni}(u, \{c_i, 1 \leq i \leq k\})$.*

Proof Let H_i^+ be the closed half-space through u containing b_i and H_i^- the closed half-space through u that avoids the interior of b_i . The inclusion $b_0 \cap b(u, r) \subseteq \bigcup_{i=1}^k b_i$ implies that $(H_0^+)^{\circ} \subseteq \bigcup_{i=1}^k (H_i^+)^{\circ}$ and $\bigcap_{i=1}^k H_i^- \subseteq H_0^-$. We consider two cases:

- (a) The open half-space $(H_0^+)^{\circ}$ is equal to one of the open half-spaces $(H_i^+)^{\circ}$, say $(H_1^+)^{\circ}$. Clearly, $c_0 \in \text{coni}(u, \{c_1\})$.
- (b) The open half-space $(H_0^+)^{\circ}$ is distinct from all other half-spaces $(H_i^+)^{\circ}$ and contained in the union of two of them, say $(H_1^+)^{\circ}$ and $(H_2^+)^{\circ}$. In that case, u and the two centers c_1 and c_2 are not aligned and the vectors $c_1 - u$ and $c_2 - u$ span \mathbb{R}^2 . It follows that there exist two real numbers $\lambda_1 \neq 0$ and $\lambda_2 \neq 0$ such that $c_0 - u = \lambda_1(c_1 - u) + \lambda_2(c_2 - u)$. Let us prove that $\lambda_1 > 0$ and $\lambda_2 > 0$. Consider a vector t_1 orthogonal to $c_1 - u$ and choose it so that $u + t_1 \in H_2^-$. We have $\langle t_1, c_0 - u \rangle = \lambda_2 \langle t_1, c_2 - u \rangle$. By choice of t_1 , we have $u + t_1 \in H_1^- \cap H_2^- \subseteq H_0^-$ and therefore $\langle t_1, c_0 - u \rangle \leq 0$ and $\langle t_1, c_2 - u \rangle \leq 0$, showing that $\lambda_2 > 0$. Similarly we get $\lambda_1 > 0$. Hence, $c_0 \in \text{coni}(u, \{c_1, c_2\})$. \square

We are now ready to give the proof of Theorem 3.

Proof of Theorem 3 Let us establish the inclusion $\text{coni}(u, C) \cap \partial b(u, \varepsilon) \subseteq w_\varepsilon(u)$. First, consider $x \in \text{coni}(u, C) \cap \partial b(u, \varepsilon)$. By Lemma 1, there exists $r > 0$ such that $b(x, \varepsilon) \cap b(u, r) \subseteq \bigcup \mathcal{B} \subseteq S$. Hence, $x \in w_\varepsilon(u)$. Second, suppose that there exists $c \in C$ such that $\{x\} = [uc] \cap \partial b(u, \varepsilon)$. Clearly, $b(x, \varepsilon) \subseteq b(c, \|c - u\|) \subseteq S$. Once again, $x \in w_\varepsilon(u)$.

Conversely, let us establish the inclusion $w_\varepsilon(u) \subseteq \text{coni}(u, C) \cap \partial b(u, \varepsilon)$. Consider $x \in w_\varepsilon(u)$. By definition, $x \in \partial b(u, \varepsilon)$ and there exists $r > 0$ such that $b(x, \varepsilon) \cap b(u, r) \subseteq S$. Since \mathcal{B} represents S locally around u , we also have $b(x, \varepsilon) \cap b(u, r) \subseteq \bigcup \mathcal{B}$ for r small enough. It follows by Lemma 2 that $x \in \text{coni}(u, C) \cap \partial b(u, \varepsilon)$. If $x \in \text{coni}(u, C)$, we are done. Otherwise, x lies on the boundary of $\text{coni}(u, C)$. This boundary consists of half-lines originating from u and going through some point of C . Let $c \in C$ such that c supports one such half-line and $x \in [uc]$. Without loss of generality, we may assume that the half-line $[uc]$ only meets C at c itself. Let us prove that $x \in [uc]$. Consider b the disk of \mathcal{B} centered at c , and notice that, since $[uc] \subseteq \partial \text{coni}(u, C)$, b supports an edge of ∂S in the neighbourhood of u . Assume by contradiction that $x \in [uc] \setminus [uc]$. Then $b \subseteq b(x, \varepsilon)$. In particular $b(x, \varepsilon)$ contains in its interior and in a neighborhood of u part of the edge supported by b , implying that $x \notin w_\varepsilon(u)$. \square

3 Rational (δ, ε) -ball approximation is in NP

To check whether a finite collection of disks \mathcal{B} is a (δ, ε) -ball approximation of a finite union of disks $\bigcup \mathcal{S}$, we have to test the two inclusions $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq$

$\bigcup \mathcal{B} \subseteq (\bigcup \mathcal{S})^{\oplus \delta}$ in polynomial time. To do so, we build on results of Section 2 and describe a method that given the triplet $(\mathcal{S}, \varepsilon, \mathcal{B})$ returns true if and only if the first inclusion $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$ holds. To test the second inclusion, it suffices to apply the same method to the triplet $(\mathcal{B}, 0, \{s^{\oplus \delta} \mid s \in \mathcal{S}\})$.

Before describing our method in Section 3.2, we explain in Section 3.1 how to encode the erosion (or union) of a finite union of disks using cells of an arrangement of circles.

3.1 Capturing the erosion through an arrangement

We recall that a finite set of circles \mathcal{C} induces a subdivision of the plane that consists of vertices, edges and faces. A vertex is an intersection point of two circles in \mathcal{C} . An edge is a maximal connected portion of a circle that is not intersected by any other circle in \mathcal{C} . A face is a maximal connected region of \mathbb{R}^2 not intersected by any circle in \mathcal{C} . The subdivision is referred to as the arrangement induced by \mathcal{C} . Vertices, edges and faces are called the cells of the arrangement. Arrangements induced by circles have a size quadratic in the number of circles and can be computed in polynomial time [29].

We say that a set of cells in an arrangement *captures* a set $X \subseteq \mathbb{R}^2$ if we can label cells as IN or OUT in such a way that cells with label IN form a partition of the set X . In this section, we show that the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$ can be captured by the cells of the arrangement induced by the set of circles:

$$\mathcal{C}_\varepsilon(\mathcal{S}) = \{\partial b(v, \varepsilon) \mid v \in \mathcal{V}\} \cup \{\partial b^{\ominus \varepsilon} \mid b \in \mathcal{S} \text{ and } \text{radius}(b) \geq \varepsilon\}.$$

We shall see that labeling the arrangement can be done in polynomial time. Furthermore, the arrangement can be computed in an exact manner [35] because all circles involved in the construction are rational (center coordinates and radii are rational numbers).

Denote by $\mathcal{A}_\varepsilon(\mathcal{S})$ the arrangement formed by circles in $\mathcal{C}_\varepsilon(\mathcal{S})$. We already know from Section 2.3 that the wavefront is composed of circular arcs supported by the set of circles $\mathcal{C}_\varepsilon(\mathcal{S})$. The boundary of the erosion $\partial S^{\ominus \varepsilon}$ being a subset of the wavefront, it is also composed of circular arcs supported by $\mathcal{C}_\varepsilon(\mathcal{S})$. It is easy to check that the boundary of the erosion is captured by the vertices and edges of $\mathcal{A}_\varepsilon(\mathcal{S})$ (see Figure 2.c). It follows that every cell of $\mathcal{A}_\varepsilon(\mathcal{S})$ is either inside or outside the erosion. We establish the following result:

Lemma 3 *We can label in polynomial time $\mathcal{A}_\varepsilon(\mathcal{S})$ in such a way that cells with label IN form a partition of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$.*

To label cells in the arrangement, we need the following technical lemma. We call each edge reduced to a point *degenerate*.

Lemma 4 *Let e be a non-degenerate edge of $\mathcal{A}_\varepsilon(\mathcal{S})$ that supports the boundary of the erosion. Then e is adjacent to a face in the erosion and a face in its complement.*

Proof Let $S = \bigcup \mathcal{S}$ and $x \in e$. By Theorem 2, there is $u \in \partial S$ such that $x \in w_\varepsilon(u) \setminus (\mathcal{V}^{\oplus \varepsilon})^\circ$. Let z_t be the point on half-line $[ux)$ such that $d(z_t, u) = t$ (see Figure 3.a). To establish the lemma, we proceed in three steps:

(1) First, we prove that $z_t \notin (\mathcal{V}^{\oplus t})^\circ$ for all t close enough to ε . Recall that $x \notin (\mathcal{V}^{\oplus \varepsilon})^\circ$ or equivalently $d(x, \mathcal{V}) \geq \varepsilon$. We claim that if we exclude u from the vertex set \mathcal{V} , then the distance from x to the remaining vertices of \mathcal{V} can only be larger, that is, $d(x, \mathcal{V} \setminus \{u\}) > \varepsilon$. Indeed, if this is not the case, then x would be equidistant to u and a vertex $v \in \mathcal{V} \setminus \{u\}$. This would imply that x lies on at least two circles of the arrangement (the one supporting e and another bounding $b(v, \varepsilon)$). Hence, x would be a vertex of the arrangement, contradicting our assumption that x lies on edge e . It follows that we can find $T > \varepsilon$ such that for all $0 \leq t \leq T$, we have $d(z_t, \mathcal{V} \setminus \{u\}) \geq d(z_T, \mathcal{V} \setminus \{u\}) - d(z_t, z_T) > T - (T - t) = t$. Hence, $d(z_t, \mathcal{V}) \geq t$ for all $0 \leq t \leq T$ or equivalently $z_t \notin (\mathcal{V}^{\oplus t})^\circ$ for all $0 \leq t \leq T$.

(2) Second, we prove that $z_t \in w_t(u)$ for all t close enough to ε . Let F a collection of disks that represents S locally around u and C the collection of their centers. Recall that by Theorem 3, $w_t(u) = \text{coni}^*(u, C) \cap \partial b(u, t)$, hence $x \in \text{coni}^*(u, C)$. We want to prove that we also have $z_t \in \text{coni}^*(u, C)$, for t in an open neighbourhood of ε . If $x \in \text{coni}^*(u, C)$, then for all $t > 0$, $z_t \in \text{coni}^*(u, C)$. If $x \notin \text{coni}^*(u, C)$, then there is $c \in C$ such that $x \in [uc]$. Without loss of generality we can assume that $\text{coni}^*(u, C) \cap [ux) = [uc]$. We claim that $x \neq c$. Indeed let b be the disk in F centered at c and let r be its radius. Assume for a contradiction that $x = c$. Then, $r = \varepsilon$ and $b^{\ominus \varepsilon} = \{c\}$, showing that $c = x$ is a degenerate edge of the arrangement and contradicting our assumption that $x \in e$ lies on a non-degenerate edge of the arrangement. Thus, $x \neq c$ and $r > \varepsilon$. Also for all $0 \leq t \leq r$, we have $z_t \in [uc] \subseteq \text{coni}^*(u, C)$, hence $z_t \in w_t(u)$.

(3) Steps (1) and (2) imply that for t close enough to ε , we have $z_t \in w_t(u) \setminus (\mathcal{V}^{\oplus t})^\circ$. By Theorem 2, this implies that $z_t \in \partial(S^{\ominus t})$ and thus that $d(z_t, \partial S) = t$. In particular, for $t < \varepsilon$, $d(z_t, \partial S) < \varepsilon$ and $z_t \in (S^{\ominus \varepsilon})^c$ while for $t > \varepsilon$, $d(z_t, \partial S) > \varepsilon$ and $z_t \in S^{\ominus \varepsilon}$, which suffices to conclude. \square

We now prove Lemma 3.

Proof of Lemma 3 We label cells in $\mathcal{A}_\varepsilon(\mathcal{S})$ as follows:

- (a) We label as BOUNDARY all cells on the boundary of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$, using results in Section 2. More precisely, we compute first vertices and edges on the boundary of $\bigcup \mathcal{S}$ (building for instance the arrangement $\mathcal{A}_0(\mathcal{S})$). Then, for each element on the boundary of $\bigcup \mathcal{S}$, we compute its contribution to the wavefront (using Theorem 3). Each contribution is composed of cells and we label the cells belonging to the boundary of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$ as BOUNDARY. To detect such cells, we use Theorem 2 and keep those lying outside $(\mathcal{V}^{\oplus \varepsilon})^\circ$.
- (b) We label as IN each face covered by a disk $b^{\ominus \varepsilon}$ for $b \in \mathcal{S}$ and label as OUT each face covered by a disk $b(v, \varepsilon)$ for $v \in \mathcal{V}$. If no faces have received a label so far, this means that $S^{\ominus \varepsilon}$ is either empty or a finite collection of singletons.

In this case, each singleton is labeled IN, all other cells are labeled OUT and we are done. Otherwise, we continue.

(c) Starting from a face either with label IN or label OUT, we do a traversal of the arrangement $\mathcal{A}_\varepsilon(\mathcal{S})$, stepping from one face to an adjacent face. If we cross an edge labeled BOUNDARY, the face currently visited receives a label different from the face on the other side of the crossed edge. Otherwise, the face currently visited receives the same label as the face on the other side. Lemma 4 guarantees that after this step, all faces are properly labeled.

(d) To label remaining vertices and edges, we use the fact that the erosion is a closed set. We thus give label IN to every vertex or edge bounding a face labeled IN. Other vertices and edges receive the label OUT.

Clearly, all steps can be done in time polynomial in the size of the arrangement. \square

3.2 Checking set inclusion

We explain in this section how to check the inclusion $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$ in polynomial time. We compute the arrangement $\mathcal{A}_\varepsilon(\mathcal{S})$ and label its cells as IN or OUT depending on whether they belong to the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$ or not as explained in the previous section. Likewise, we compute a second arrangement $\mathcal{A}_0(\mathcal{B})$ and label its cells as IN or OUT depending on whether they belong to the union $\bigcup \mathcal{B}$ or not. We then overlay the two arrangements. This operation consists in computing the arrangement of the joint collection of circles, and it is possible to do so while inheriting the labels from both input arrangements [4, 34]. By inspecting the labels of the overlay, we can thus determine if every cell in the erosion is also in the union and conclude on the inclusion test. All operations can be done in polynomial time. In Appendix A, we provide an alternative method that does not use any overlay.

4 Rational (δ, ε) -ball approximation is NP-hard

In this section, we establish that the rational (δ, ε) -ball approximation problem in \mathbb{R}^n is NP-hard for $n = 2$. We prove NP-hardness through a reduction from a variant of the vertex cover problem. Recall that for a graph $G = (V, E)$, a subset $C \subseteq V$ is a vertex cover of G if every edge of E is incident to a vertex of C . Finding a minimum vertex cover is NP-hard, even when restricted to cubic planar graphs, that is, planar graphs of degree at most 3 [20]. We shall perform the reduction from this particular variant. Given a cubic planar graph $G = (V, E)$ and rational numbers $\delta > 0$ and $\varepsilon \geq 0$, we show below how to build a finite collection $\mathcal{S}(G, \delta, \varepsilon)$ of disks in the plane such that G has a vertex cover of cardinality k if and only if $\bigcup \mathcal{S}(G, \delta, \varepsilon)$ admits a (δ, ε) -ball approximation of cardinality $k + L$, where L is an integer depending only on $\mathcal{S}(G, \delta, \varepsilon)$. Note that the reduction presented below works for all values $\delta > 0$.

When $\delta = 0$, a similar reduction is presented in Appendix C that requires slightly more complicated vertex gadgets for degree-three vertices.

4.1 Reduction from vertex cover

Gadgets. We build the set of disks $\mathcal{S}(G, \delta, \varepsilon)$ from the tuple (G, δ, ε) by converting each vertex of G into a *vertex* gadget and each edge of G into an *edge* gadget. The vertex gadget for a vertex $v \in V$ consists of disks of radius ε , one for each edge incident to v . The edge gadget for an edge $e \in E$ is a finite non-empty set of disks of odd length, which all share the same radius ε ; see Figure 6.b. Disks in the gadgets are obtained by first associating to each edge $e = \{v, w\}$ a finite sequence of disks of length at least three. Disks in the sequence have radius ε . The centers of those disks are called the *rotulae* of edge e . We distinguish two types of rotulae: *Extreme* rotulae are the centers of the two extreme disks in the sequence and are denoted as c_{ve} and c_{we} . *Regular* rotulae are the centers of other disks. The rotulae of two consecutive disks in the sequence are said to be neighbors of one another. Note that extreme rotulae have only one neighbor and regular rotulae have two neighbors. The edge gadget for e is the set of disks in the sequence centered at regular rotulae. The vertex gadget for v is the set of disks $b(c_{ve}, \varepsilon)$, where e ranges over all edges incident to v . Hence, a vertex gadget may contain up to three disks depending on the degree of the vertex that generated the gadget. We say that an extreme rotulae c_{ve} and a vertex x are *linked* to one another whenever $v = x$. Finally, we define $\mathcal{S}(G, \delta, \varepsilon)$ as the collection of disks resulting from the conversion of vertices and edges into vertex gadgets and edge gadgets respectively.

At this point, we haven't specified the centers of disks yet, nor the number of rotulae per edge that we need (only that it should be an odd number). We postpone to Section 4.2 the description of how disks composing gadgets are chosen. This will be done from an orthogonal grid drawing of G ; see Figure 6 for an example. At this stage it suffices to know that we can always choose disks composing gadgets so that they fulfill the properties (i)-(viii) listed below. To simplify notations, we shall refer to $\mathcal{S}(G, \delta, \varepsilon)$ and $S(G, \delta, \varepsilon) = \bigcup \mathcal{S}(G, \delta, \varepsilon)$ simply as \mathcal{S} and S .

- (i) The ε -erosion $S^{\ominus \varepsilon}$ is exactly the collection of rotulae.
- (ii) Any disk $b \subseteq S^{\oplus \delta}$ contains at most 2 regular rotulae.
- (iii) If a disk $b \subseteq S^{\oplus \delta}$ contains 2 regular rotulae, then they are neighbors of each other and b does not contain any other rotula (neither regular nor extreme).
- (iv) Let c be a regular rotula, c_1 and c_2 its two neighbors. There exist two disks $b_1, b_2 \subseteq S^{\oplus \delta}$ such that $\{c_1, c\} \subseteq b_1$ and $\{c_2, c\} \subseteq b_2$.
- (v) Any disk $b \subseteq S^{\oplus \delta}$ contains at most 3 extreme rotulae.
- (vi) If a disk $b \subseteq S^{\oplus \delta}$ contains 2 or 3 extreme rotulae, then these extreme rotulae are linked to the same vertex and b only contains extreme rotulae linked to that vertex.

- (vii) Let v be a vertex. There exists a disk $b \subseteq S^{\oplus\delta}$ that contains all extreme rotulae linked to v .
- (viii) If a disk $b \subseteq S^{\oplus\delta}$ contains both a regular rotula and an extreme rotula, then these are neighbors and b does not contain any other rotula (neither regular nor extreme).




Fig. 6: Conversion from (a) an orthogonal grid drawing of a graph G with 5 vertices and 5 edges to (b) the collection of disks $\mathcal{S}(G, \delta, \varepsilon)$ with 10 (purple) vertex rotulae and 31 (blue) edge rotulae. The δ -dilation of the union of disks is bounded by the dashed blue lines.

Canonical coverings. We ask for the smallest number of disks contained in $S^{\oplus\delta}$ that suffices to cover the rotulae of a given edge $e = \{v, w\} \in E$. Recall that the rotulae of e include the two extreme rotulae c_{ve} and c_{we} used to build the vertex gadgets of v and w in addition to the regular rotulae used to build the edge gadget of e . By construction, e has an odd number of rotulae greater than or equal to 3. Let $2\ell(e) + 1$ be this number, where $\ell(e) \geq 1$ is an integer. From Property (ii) we need at least $\lceil (2\ell(e) - 1)/2 \rceil = \ell(e)$ disks in $S^{\oplus\delta}$ in order to cover regular rotulae of e . By (iv) there always exists a collection of $\ell(e)$ disks in $S^{\oplus\delta}$ which covers these regular rotulae plus one of the two extreme

rotulae of e , and that extreme rotula can be chosen arbitrarily. Indeed, it suffices to cover pairs of neighbouring rotulae with disks in $S^{\oplus\delta}$, while making sure that the extreme rotula to cover and its neighbour are one of these pairs (see Figure 7). This gives two possible coverings of the regular rotulae of e (one for each extreme rotulae of e) which we shall refer to as *canonical* for e . Furthermore, properties (iii) and (viii) guarantee that any disk containing a regular rotula will only contain rotulae of the same edge. Therefore it is necessary and sufficient to use $\ell(e)$ disks in $S^{\oplus\delta}$ to cover the regular rotulae of an edge e , and these $\ell(e)$ disks exclusively cover rotulae of e . However, one of the two extreme rotula of e will not be covered by these $\ell(e)$ disks and one extra disk is required to cover all rotulae of e for a total of $\ell(e) + 1$ disks. Unlike the previous $\ell(e)$ disks, by (v) this extra disk may be used to cover extreme rotulae of several edges assuming these extreme rotulae are all linked to the same vertex. We define $L = \sum_{e \in E} \ell(e)$ which is the smallest number of disks needed to cover all regular rotulae and which can be achieved by taking for each edge, one of its two possible canonical coverings.




Fig. 7: A canonical covering of an edge gadget with two green disks. Same color convention as in Figure 6.b.

Property 4 If G has a vertex cover $C \subseteq V$, then S has a (δ, ε) -ball approximation \mathcal{B} with $|\mathcal{B}| = L + |C|$.

Proof By (i), we know that any (δ, ε) -ball approximation of S is a collection of disks contained in $S^{\oplus\delta}$ whose union covers the rotulae of all edges of G . To build such a (δ, ε) -ball approximation, we use Property (vii) and start by selecting for each vertex $v \in C$ a disk that covers all extreme rotulae linked to v . This gives a set \mathcal{C} of $|C|$ disks that cover at least one extreme rotula per edge. Using the appropriate canonical covering of each edge, that is, the one excluding the extreme rotula already covered by \mathcal{C} , we complete the set \mathcal{C} into a (δ, ε) -ball approximation adding L more disks. \square

Property 5 If S has a (δ, ε) -ball approximation \mathcal{B} , then G has a vertex cover C with $|C| \leq |\mathcal{B}| - L$.

Proof Without loss of generality, we may assume that all disks in \mathcal{B} cover at least one rotula. Indeed, if a disk b does not cover any rotula, it can be removed from \mathcal{B} while keeping the property that \mathcal{B} is a (δ, ε) -ball approximation.

Starting from \mathcal{B} , we first deduce a (δ, ε) -ball approximation \mathcal{R} of S having the property that it contains one of the two canonical coverings of each edge $e \in E$. For $e \in E$, let $\mathcal{B}_e = \{b \in \mathcal{B} \mid b \text{ contains a regular rotula of } e\}$. Note that for different edges the \mathcal{B}_e are disjoint and that $|\mathcal{B}_e| \geq \ell(e)$. Given an extreme rotula c_{ve} , we denote by \mathcal{C}_{ve} the canonical covering that contains c_{ve} and all the regular rotulae of edge e . Initializing \mathcal{R} to \mathcal{B} , we modify \mathcal{R} as follows, applying the following transformation to each edge e of G :

- If disks in \mathcal{B}_e cover no extreme rotula, we replace \mathcal{B}_e by one of the two canonical coverings of e .
- If disks in \mathcal{B}_e cover 1 extreme rotula, we replace \mathcal{B}_e by the canonical covering of e that covers this extreme rotula.
- If disks in \mathcal{B}_e contains both extreme rotulae of e , then $|\mathcal{B}_e| \geq \ell(e) + 1$. We choose arbitrarily one of the two extreme rotulae, say c_{ve} , and let b be a disk in $S^{\oplus \varepsilon}$ containing the other extreme rotula c_{we} but no regular rotula. We replace \mathcal{B}_e by $\mathcal{C}_{ve} \cup \{b\}$.

Each of these substitutions preserves the (δ, ε) -ball approximation property and does not increase the cardinality of the resulting collection of disks. Consider the disks of \mathcal{R} that do not contain any regular rotula, $\mathcal{C} = \mathcal{R} \setminus (\cup_{e \in E} \mathcal{R}_e)$. By construction, $|\mathcal{C}| = |\mathcal{R}| - L \leq |\mathcal{B}| - L$ and \mathcal{C} covers at least half of the extreme rotulae (at least one for each edge). Let C be the set of vertices $v \in V$ linked to an extreme rotula covered by a disk in \mathcal{C} . We claim that C is a vertex cover of G and that its cardinality satisfies $|C| \leq |\mathcal{C}|$. All $b \in \mathcal{C}$ must contain at least one extreme rotula, thus C is empty if and only if \mathcal{C} is empty. In this particular case, G has no edges because otherwise \mathcal{R} would only cover half of the extreme rotulae. It follows that in this case the empty set is a vertex cover of G . Assume now that \mathcal{C} is not empty. By construction, any $b \in \mathcal{C}$ covers only extreme rotulae, which by (vi) are all linked to the same vertex v . It is thus possible to map each $b \in \mathcal{C}$ to one vertex $v \in C$, showing that $|C| \leq |\mathcal{C}|$. To see that C is a vertex cover, recall that \mathcal{C} covers exactly one of the two extreme rotulae of each edge. The definition of C thus ensures it contains at least one endpoint of each edge. \square

4.2 Practical construction

Given as input a graph G and rational numbers $\delta > 0$ and $\varepsilon \geq 0$, we build a collection of rational disks fulfilling properties (i) to (viii). To do so, we first relax the condition that the collection is rational and build a collection of disks \mathcal{S} that satisfies properties (i) to (viii). We then show at the end of the section how to perturbate \mathcal{S} so as to make disks in the collection rational while maintaining properties (i) to (viii). To build \mathcal{S} , we rely on the following result:

Theorem 4 ([9]) *There is a linear time and space algorithm to draw a connected cubic graph on an orthogonal grid.*

Recall that the drawing of a graph G on an orthogonal grid is a drawing of G such that each vertex is drawn as a vertex grid and each edge is drawn as a chain of horizontal and vertical segments of the grid; see Figure 6.a for an example of such a drawing. Given a drawing of G on an orthogonal grid, we describe a way to convert it into an appropriate collection of disks \mathcal{S} . To perform the conversion, we fix the size of the grid to $8(\delta + \varepsilon)$ so that we can fit square blocks of size $4(\delta + \varepsilon) \times 4(\delta + \varepsilon)$ as in Figure 8. Note that in Figures 6 to 10, we set $\varepsilon = 2\delta$.




Fig. 8: Drawing (in blue) of a cubic graph on an orthogonal grid and grid division into blocks. Gray lines represent the grid and dashed red lines are the blocks.

There are two different ways in which the blocks meet the graph drawing: the block either contains one vertex or only a portion of one edge. Blocks containing a vertex only vary depending on the number and layout of incident edges. Similarly, blocks containing a portion of an edge vary depending on whether the edge bends within the block or not. In each case, we convert the graph drawing covered by the block into a set of disks of radii ε . For blocks containing a vertex, see Figure 9 for the four subcases. Similarly for edges, we have two subcases. However, recall that edge gadgets must have an odd number of rotulae. To achieve this, we use the fact that every edge necessarily crosses at least one block in a straight line and provide an odd and an even conversion for this type of block. The three block conversions are presented in Figure 10. Finally, we build \mathcal{S} by converting the orthogonal graph drawing of G into a set of disks using the 7 types of blocks described in Figures 9 and 10. Each block contains between 2 and 6 rotulae, the positions of which depend on parameters δ and ε . An extra parameter t is also used within two blocks.

Theorem 5 *For all values $1 < t < 2 - \frac{1}{\sqrt{2}}$, the collection of disks \mathcal{S} resulting from the conversion of the orthogonal graph drawing satisfies properties (i) to (viii).*

Before proving the theorem, we give some definitions and study the medial axes of $S = \bigcup \mathcal{S}$ and $S^{\oplus \delta}$. For simplicity, we write b_x for $b(x, \varepsilon)$ and B_x for $b(x, \delta + \varepsilon)$. We also identify a vertex $v \in V$ and its drawing as a vertex

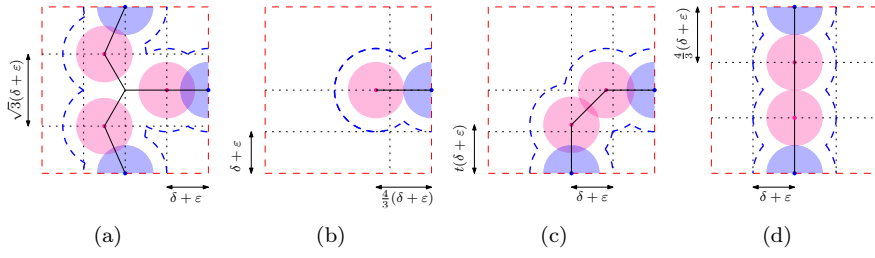


Fig. 9: Block conversions for (a) a degree-3 vertex; (b) a degree-1 vertex; (c) a degree-2 vertex in a bend; (d) a degree-2 vertex in a straight-line. The red dashed square delimits the block. The medial axis of $S^{\oplus\delta}$ that intersects each block is the set of black segments. Same color convention as in Figure 6.b.




Fig. 10: Block conversions (a) when the edge bends within the block; (b) and (c) when the edge traverses the block in a straight line. (b) uses an even number of disks and (c) an odd number. The medial axis of $S^{\oplus\delta}$ that intersects each block is the set of black segments. Same color convention as in Figure 6.b.

grid. Given two disks $b_1 = b(c_1, r_1)$ and $b_2 = b(c_2, r_2)$ we recall that a *convex combination* of b_1 and b_2 is a disk whose equation is $\alpha\pi_1(x) + (1 - \alpha)\pi_2(x) \leq 0$ where $\alpha \geq 0$ and $\pi_i(x) = \|x - c_i\|^2 - r_i^2$ for $1 \leq i \leq 2$. If b is the convex combination of b_1 and b_2 determined by coefficient α , then the center of b is $\alpha c_1 + (1 - \alpha)c_2$ and ∂b passes through $\partial b_1 \cap \partial b_2$. Recall that disks in \mathcal{B} are *independent* if for every subset $\mathcal{B}' \subseteq \mathcal{B}$, including $\mathcal{B}' = \emptyset$, there is a point in $\bigcap \mathcal{B}'$ that does not belong to any disk not in \mathcal{B}' . If three disks b_1, b_2 and b_3 are independent, then the boundary of their union $\bigcup_{1 \leq i \leq 3} b_i$ has exactly three vertices and the circumdisk of these vertices is contained in the union of the three disks. We call the circumdisk the *branching disk* of the collection $\{b_1, b_2, b_3\}$. Because of the way in which rotulae are placed within each block, the following property holds:

Property 6 Consider a degree-three vertex v and let c_1, c_2 and c_3 be the three extreme rotulae linked to v . The three disks b_{c_1}, b_{c_2} and b_{c_3} may intersect pairwise but do not have a common intersection. The three disks B_{c_1}, B_{c_2} and B_{c_3} are independent.

Let us denote by B_v^* the branching disk of the collection $\{B_{c_1}, B_{c_2}, B_{c_3}\}$. Observe that since we place the three extreme rotulae c_1, c_2 and c_3 linked to

v in such a way that they form an equilateral triangle, the circumdisk of c_1 , c_2 and c_3 is B_v and furthermore $B_v^* = B_v$.

Property 7 For $1 < t < 2 - \frac{1}{\sqrt{2}}$, medial disks of $S = \bigcup \mathcal{S}$ are

- disks of \mathcal{S} ;
- convex combinations of disks b_{c_1} and b_{c_2} such that $\partial b_{c_1} \cap \partial b_{c_2}$ consists of two points and one of the following holds:
 - (1) c_1 and c_2 are two neighboring rotulae;
 - (2) c_1 and c_2 are extreme rotulae linked to the same vertex.

Property 8 For $1 < t < 2 - \frac{1}{\sqrt{2}}$, medial disks of $S^{\oplus\delta}$ are

- disks centered at rotulae with radius $\delta + \varepsilon$;
- convex combinations of disks B_{c_1} and B_{c_2} where c_1 and c_2 are two neighboring rotulae;
- convex combinations of disks B_{c_1} and B_{c_2} where c_1 and c_2 are extreme rotulae linked to the same degree-two vertex;
- convex combinations of disks B_v^* and B_c where v is a degree-three vertex and c is an extreme rotulae linked to v .

We are now ready to prove Theorem 5.

Proof of Theorem 5 Property (i) holds because Property 7 implies that every medial disk b of S is the convex combination of two (possibly identical) disks of \mathcal{S} . Since all disks of \mathcal{S} share the same radius ε , the radius of b does not exceed ε and is equal to ε if and only if $b \in \mathcal{S}$. It follows that the ε -erosion of S is the collection of rotulae.

Let us prove that every disk B of $S^{\oplus\delta}$ contains at most two rotulae, except medial disks B_v centered at degree-three vertices v which contain the three extreme rotulae linked to v . Notice that it suffices to establish this for medial disks of $S^{\oplus\delta}$ since any disk $B \subseteq S^{\oplus\delta}$ is always contained in a medial disk of $S^{\oplus\delta}$. Consider a medial disk B of $S^{\oplus\delta}$. Property 8 implies that one of the following two situations occurs:

- B is the convex combination of disks B_{c_1} and B_{c_2} where c_1 and c_2 are two neighboring rotulae or two extreme rotulae linked to the same degree-two vertex. In that case, it is easy to check that the only rotulae contained in $B_{c_1} \cup B_{c_2}$ are c_1 and c_2 and since $B \subseteq B_{c_1} \cup B_{c_2}$, the only rotulae that B may contain are c_1 and c_2 .
- B is the convex combination of disks B_v^* and B_c where v is a degree-three vertex and c is an extreme rotula linked to v . In that case, $B_v^* \cup B_c$ contains four rotulae; see Figure 11. One of them is c , two of them lie on $\partial B_v^* \setminus B_c$ and one of them lie on $\partial B_c \setminus B_v^*$. It follows that if $B \neq B_v^*$, then B avoids $\partial B_v^* \setminus B_c$ while being contained in $B_v^* \cup B_c$, implying that B contains at most two rotulae. If $B = B_v^*$, then B contains the three extreme rotulae linked to v .




Fig. 11: $B_v^* \cup B_c$ contains four rotulae : three extreme rotulae in pink, and a regular one in blue.

From there, by carefully studying the type of rotulae contained in disks of $S^{\oplus\delta}$, Properties (ii), (iii), (v), (vi), (viii) follow easily.

Property (iv) holds because the distance between two neighboring rotulae c_1 and c_2 has been adjusted so that $\|c_1 - c_2\| \leq \sqrt{2}(\delta + \varepsilon)$, implying that the convex combination of B_{c_1} and B_{c_2} centered at the midpoint of segment c_1c_2 contains both c_1 and c_2 ; see Figure 7.

Property (vii) holds because one can always find a disk of $S^{\oplus\delta}$ that contains all extreme rotulae c_1, \dots, c_k linked to a degree- k vertex v . If $k = 1$, take B_{c_1} ; if $k = 2$, take the convex combination of B_{c_1} and B_{c_2} centered at the midpoint of segment c_1c_2 ; if $k = 3$, take B_v . \square

We now explain how to perturbate the set \mathcal{S} so as to get a collection of rational disks which fulfills properties (i) to (viii). Observe that by construction we only need to perturbate coordinates of some of the disks centered at extreme rotulae linked to degree-three vertices. Precisely, let v be a degree-three vertex whose extreme rotulae are c_1, c_2 and c_3 . Assume c_1 and c_2 lie on the same vertical line L . Then, the coordinates of the three points are rational except the y -coordinates of c_1 and c_2 . We perturbate c_1 and c_2 so as to get rational points by moving c_1 and c_2 along the vertical line L towards each other so as to decrease $\|c_1 - c_2\|$ while keeping the property that $c_1c_2c_3$ is isosceles. We note that under a sufficiently small motion of points c_1 and c_2 , Properties 6 to 8 still hold. Doing this for all degree-three vertices, we get a collection of rational disks which still satisfies properties (i) to (viii). To see this, it suffices to adapt the proof of Theorem 5, showing that every disk of $S^{\oplus\delta}$ contains at most two rotulae, except some of the disks centered “near” degree-three vertices v which contain the three extreme rotulae linked to v .

5 Conclusion

The main contribution of this paper is the proof that the rational (δ, ε) -ball approximation problem is NP-complete in \mathbb{R}^2 . As a by-product, we immedi-

ately get that the problem is also NP-hard in \mathbb{R}^n for $n > 2$. Indeed, each instance \mathcal{S}_2 of the problem in \mathbb{R}^2 can be transformed in polynomial time into an instance \mathcal{S}_n of the problem in \mathbb{R}^n , replacing disks in \mathcal{S}_2 by n -dimensional balls in \mathbb{R}^n with the same centers and radii. It is easy to see that \mathcal{S}_2 has a (δ, ε) -ball approximation of cardinality k if and only if \mathcal{S}_n has a (δ, ε) -ball approximation of cardinality k .

In our proof of NP-hardness in \mathbb{R}^2 , we have chosen carefully the set of disks composing the shape, so that its erosion is composed of a finite set of points (the rotulae). It follows that the reduction from vertex cover induces a set cover problem of a finite set of points (the rotulae), using an infinite range of balls (the medial axis of the dilation). We notice however that it is possible to define a similar construction where the erosion of the shape is an infinite set of points too. Furthermore, a construction where the shape, its erosion and its dilation are arcwise-connected is also possible. To build such variants, it suffices to tune the radii of disks composing the gadgets, and distances between consecutive disks so that the eight properties (i)-(viii) are fulfilled. This does not change the overall result.

Ascertaining the computational complexity of a problem is often the first step before the design of algorithms that return either an optimal solution for a simple subclass of inputs, or an approximated solution for wider classes of inputs. In [28], we investigate the (δ, ε) -ball approximation problem for the subclass of shapes S defined as a finite union of balls in \mathbb{R}^2 , with an acyclic medial axis $\text{MA}(S)$. We design a greedy algorithm that computes an optimal (δ, ε) -ball approximation of S in polynomial time with respect to the size of $\text{MA}(S)$. Whether or not this algorithm can be extended for acyclic finite unions of balls in \mathbb{R}^3 remains an open question.

In the (δ, ε) -ball approximation definition, the core set to cover and the outer set that must not be crossed over are defined through mathematical morphology operators. It is natural to investigate an even more general version of the problem:

Problem 3 Given a shape $S \subseteq \mathbb{R}^n$, a subset S^- of S and a superset S^+ of S , find the smallest collection of balls whose union X satisfies $S^- \subseteq X \subseteq S^+$.

When the shape is defined as a finite union of balls $S = \bigcup \mathcal{S}$, an obvious choice for S^- is to consider the union of balls obtained by deflating balls in \mathcal{S} , either decreasing each radius by a constant ε or multiplying each radius by a factor $\mu < 1$ as in [27]. In both cases, keeping for S^+ the δ -dilation, it is not difficult to see that the associated decision problem is again NP-complete, using the same reduction. We ask what would be the complexity of Problem 3 for other choices of S^- and S^+ .

References

1. Agarwal, P.K., Pan, J.: Near-linear algorithms for geometric hitting sets and set covers. In: Proceedings of the thirtieth annual symposium on Computational geometry, p. 271. ACM (2014)

2. Amenta, N., Kolluri, R.K.: Accurate and efficient unions of balls. In: Proceedings of the Sixteenth Annual Symposium on Computational Geometry, SCG '00, pp. 119–128. ACM (2000)
3. Attali, D., Boissonnat, J.D., Edelsbrunner, H.: Stability and computation of medial axes - a state-of-the-art report. In: T. Möller, B. Hamann, R.D. Russell (eds.) *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, pp. 109–125. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
4. Berberich, E., Eigenwillig, A., Hemmer, M., Hert, S., Mehlhorn, K., Schömer, E.: A computational basis for conic arcs and boolean operations on conic polygons. In: R. Möhring, R. Raman (eds.) *Algorithms — ESA 2002*, pp. 174–186. Springer Berlin Heidelberg (2002)
5. Bradshaw, G., O’Sullivan, C.: Adaptive medial-axis approximation for sphere-tree construction. *ACM Transactions on Graphics (TOG)* **23**(1), 1–26 (2004)
6. Brönnimann, H., Goodrich, M.T.: Almost optimal set covers in finite vc-dimension. *Discrete & Computational Geometry* **14**(4), 463–479 (1995)
7. Broutta, A., Coeurjolly, D., Sivignon, I.: Hierarchical discrete medial axis for sphere-tree construction. In: *International Workshop on Combinatorial Image Analysis*, pp. 56–67. Springer (2009)
8. Cabello, S., De Berg, M., Giannopoulos, P., Knauer, C., Van Oostrum, R., Veltkamp, R.C.: Maximizing the area of overlap of two unions of disks under rigid motion. *International Journal of Computational Geometry and Applications* **19**(06), 533–556 (2009)
9. Calamoneri, T., Petreschi, R.: An efficient orthogonal grid drawing algorithm for cubic graphs. In: *Computing and Combinatorics*, pp. 31–40. Springer (1995)
10. Cazals, F., Dreyfus, T., Sachdeva, S., Shah, N.: Greedy geometric algorithms for collection of balls, with applications to geometric approximation and molecular coarse-graining. In: *Computer Graphics Forum*, vol. 33, pp. 1–17. Wiley Online Lib. (2014)
11. Chaussard, J., Couprie, M., Talbot, H.: Robust skeletonization using the discrete λ -medial axis. *Pattern Recognition Letters* **32**(9), 1384 – 1394 (2011)
12. Chazal, F., Lieutier, A.: The λ -medial axis. *Graphical Models* **67**(4), 304 – 331 (2005)
13. Chvatal, V.: A greedy heuristic for the set-covering problem. *Mathematics of operations research* **4**(3), 233–235 (1979)
14. Coeurjolly, D., Montanvert, A.: Optimal separable algorithms to compute the reverse euclidean distance transformation and discrete medial axis in arbitrary dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**(3), 437–448 (2007)
15. Culver, T., Keyser, J., Manocha, D.: Accurate computation of the medial axis of a polyhedron. In: *Proceedings of the Fifth ACM Symposium on Solid Modeling and Applications, SMA '99*, pp. 179–190. ACM (1999)
16. Dey, T.K., Zhao, W.: Approximate medial axis as a voronoi subcomplex. *Computer-Aided Design* **36**(2), 195 – 202 (2004)
17. Edelsbrunner, H., Koehl, P.: The geometry of biomolecular solvation. *Combinatorial and computational geometry* **52**, 243–275 (2005)
18. Federer, H.: Curvature measures. *Trans. Amer. Math. Soc.* **93**, 418–491 (1959)
19. Feinauer, J., Spettil, A., Manke, I., Strege, S., Kwade, A., Pott, A., Schmidt, V.: Structural characterization of particle systems using spherical harmonics. *Materials Characterization* **106**, 123 – 133 (2015)
20. Garey, M.R., Johnson, D.S.: The rectilinear steiner tree problem is NP-complete. *SIAM Journal on Applied Mathematics* **32**(4), 826–834 (1977)
21. Har-Peled, S.: *Geometric approximation algorithms*, vol. 173. American mathematical society Boston (2011)
22. Hatcher, A.: *Algebraic topology*. Cambridge University Press (2002)
23. Hubbard, P.M.: Collision detection for interactive graphics applications. *IEEE Transactions on Visualization and Computer Graphics* **1**(3), 218–230 (1995)
24. Hubbard, P.M.: Approximating polyhedra with spheres for time-critical collision detection. *ACM Transactions on Graphics (TOG)* **15**(3), 179–210 (1996)
25. Johnson, D.S.: Approximation algorithms for combinatorial problems. *Journal of computer and system sciences* **9**(3), 256–278 (1974)
26. Kershner, R.: The number of circles covering a set. *American Journal of Mathematics* **61**(3), 665–671 (1939)

27. Miklos, B., Giesen, J., Pauly, M.: Discrete scale axis representations for 3d geometry. In: *ACM Transactions on Graphics (TOG)*, vol. 29, p. 101. ACM (2010)
28. Nguyen, T.B., Sivignon, I.: Epsilon-covering: a greedy optimal algorithm for simple shapes. In: *Canadian Conference on Computational Geometry*, pp. 187–194. Vancouver (2016)
29. Pach, J.: *New Trends in Discrete and Computational Geometry. Algorithms and Combinatorics*. Springer Berlin Heidelberg (2012)
30. Ranjan, V., Fournier, A.: Matching and interpolation of shapes using unions of circles. *Computer Graphics Forum* **15**(3), 129–142 (1996)
31. de Rezende, P.J., Miyazawa, F.K., Sasaki, A.T.: A {PTAS} for the disk cover problem of geometric objects. *Operations Research Letters* **41**(5), 552 – 555 (2013)
32. Serra, J.: *Image analysis and mathematical morphology*, v. 1. Academic press (1982)
33. Tóth, G.F.: *Handbook of discrete and computational geometry*. chap. Packing and covering. Chapman & Hall, CRC (2004)
34. Wein, R., Berberich, E., Fogel, E., Halperin, D., Hemmer, M., Salzman, O., Zukerman, B.: *2D Arrangements*. CGAL: Computational Geometry Algorithms Library
35. Wein, R., Zukerman, B.: Exact and efficient construction of planar arrangements of circular arcs and line segments with applications. Tech. rep., Tech. Rep. ACS-TR-121200-01, Tel Aviv University, Israel (2006)

A An alternative way to check set inclusion

In this appendix, we provide an alternative polynomial time algorithm to check whether the erosion of a finite collection of disks is contained into the union of a finite collection of disks (see Section A.3). The algorithm relies on a characterization of the inclusion of one set into another. We give first a very general formulation (Lemma 5) before proposing a specialization to our case (Lemma 6).

A.1 A very general result

We start by proving a very general result to characterize the inclusion of one set into another one. We assume very little on the two sets, only that they are subsets of a connected topological space. We show that it suffices to test “local” inclusions at the boundary points of the two sets.

Lemma 5 *Let A and $X \neq \emptyset$ be two subsets of a connected topological space. If*

- (i) $\partial A \subseteq \bar{X}$
- (ii) $\forall x \in \partial X, \exists N_x$ an open neighbourhood of x such that $N_x \cap A \subseteq X$

then $A \subseteq X$.

Proof Letting $H = A \cap X^c$, we prove the lemma in two stages:

(a) First, we prove that $\partial H \subseteq \partial X$. Using $\partial(A \cap B) \subseteq (\partial A \cap \bar{B}) \cup (\bar{A} \cap \partial B)$ and plugging $B = X^c$, we get $\partial H \subseteq (\partial A \cap \bar{X}^c) \cup (\bar{A} \cap \partial X^c)$. Since $\partial A \subseteq \bar{X}$ and $\partial X^c = \partial X$, we deduce $\partial H \subseteq \partial X$.

(b) Second, we prove that $H = \emptyset$. Suppose for a contradiction that $H \neq \emptyset$. Since X is not empty, X^c is not the entire topological space, so that H cannot be the entire topological space neither. Thus, $\partial H \neq \emptyset$ and we can pick a point $x \in \partial H$. Because $\partial H \subseteq \partial X$, we can apply (ii) on point x , showing that there is N_x a neighbourhood of x such that $N_x \cap A \subseteq X$. Since $x \in \partial H$, necessarily $N_x \cap H \neq \emptyset$. However, $N_x \cap H = N_x \cap A \cap X^c \subseteq X \cap X^c = \emptyset$, yielding a contradiction. \square

A.2 Specialization to finite union of disks

We now give a specialized version of Lemma 5 in which one of the set is a finite union of disks.

Lemma 6 *Let A be a bounded subset of \mathbb{R}^2 and consider $X = \bigcup \mathcal{X}$ a finite union of disks.*

Let \mathcal{V} be the collection of vertices of ∂X . If

- (i) $\partial A \subseteq X$
- (ii) $\forall v \in \mathcal{V}, \exists N_v$ an open neighborhood of v such that $N_v \cap A \subseteq X$

then $A \subseteq X$.

Recall that a *hole* of a subset is a bounded connected component of its complement. The proof of the previous lemma relies on the following technical lemma.

Lemma 7 *Let H be a hole of a finite union of disks S , and \mathcal{V} the collection of vertices of ∂S . We have $\partial H \cap \mathcal{V} \neq \emptyset$.*

Proof Let \mathcal{S} be a collection of disks such that $S = \bigcup \mathcal{S}$. Without loss of generality, we assume that $\forall b \in \mathcal{S}, \partial b \cap \partial H \neq \emptyset$. Indeed, by removing any disk that does not contribute to the boundary of H , the resulting union of disks still admits H as a hole. Notice that $\partial H \subseteq \partial S$. In fact, the boundary of H is formed by vertices and edges of ∂S . By contradiction, assume there are no vertices in ∂H . Then ∂H must be an union of complete circular arcs, thus disks in \mathcal{S} cannot intersect one another. S is an union of disjoint disks and S^c has only one connected component by Alexander duality Theorem [22]. This connected component is unbounded, contradicting the existence of H . \square

Proof of Lemma 6 Let $H = A \cap X^c$. The scheme of the proof is similar to the proof of Lemma 5, with only one extra stage in the middle:

(a) First, we prove that $\partial H \subseteq \partial X$. For this, we proceed as in the first stage of the proof of Lemma 5.

(b) Second, we prove that H is either empty or an union of holes of X . The proof is given below.

(c) Last, we prove that $H = \emptyset$. For this, we proceed as in the last stage of the proof of Lemma 5, with the only difference being that we pick a point x on ∂H among vertices of ∂X . This is always possible thanks to Lemma 7 (stated above).

It remains to prove that H is either empty or an union of holes of X . Suppose $H \neq \emptyset$ and let $p \in H$. Consider H_p the connected component of H that contains p . Likewise, let X_p^c be the connected component of X^c that contains p . We prove in the following paragraph that $H_p = X_p^c$. Since A is bounded, H_p is bounded too and H_p is a hole of X , so that H is an union of holes.

Let us prove that $H_p = X_p^c$. Recall that $H_p \subseteq X^c$ is connected by definition, and that it intersects the connected component X_p^c of X^c . Hence, $H_p \subseteq X_p^c$. Let us prove the converse inclusion $X_p^c \subseteq H_p$. Let $q \in X_p^c$ and consider $\gamma : [0, 1] \rightarrow \mathbb{R}^2$ a path in X_p^c connecting p to q . Such a path exists because X_p^c is open and connected.

We have $\partial H_p \subseteq \partial H$, hence $\partial H_p \subseteq \partial X$. We have $\gamma \cap \partial H_p \subseteq \gamma \cap \partial X = \emptyset$. We claim that $q \in H_p$. Indeed, suppose for a contradiction that γ connects $p \in H_p$ to $q \notin H_p$. Then, γ would meet ∂H_p , which is impossible. \square

A.3 A polynomial time algorithm

We are now ready to describe an algorithm that tests in polynomial time whether $A \subseteq X$, where A is the erosion of a finite collection of disks and X is the union of a finite collection of disks. The algorithm hinges on Lemma 6 and its pseudocode is given in Algorithm 1. To describe it, let \mathcal{S} and \mathcal{B} be two finite collections of disks, \mathcal{C} a finite collection of circular arcs, x a point, $\varepsilon \geq 0$ a non-negative real number. We introduce four functions:

BOUNDARY takes as input the pair $(\mathcal{S}, \varepsilon)$ and returns the set of circular arcs that compose the boundary of the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$.

POINTLOCATION takes as input the triplet $(x, \mathcal{S}, \varepsilon)$ and returns **inner** if $x \in \overset{\circ}{A}$, **boundary** if $x \in \partial A$ and **outer** if $x \notin A$, where $A = (\bigcup \mathcal{S})^{\ominus \varepsilon}$.

INCLUSION takes as input $(\mathcal{C}, \mathcal{B})$ and returns **true** iff $(\bigcup \mathcal{C}) \subseteq (\bigcup \mathcal{B})$.

LOCALINCLUSION takes as input $(v, \mathcal{S}, \varepsilon, \mathcal{B})$ where v lies on the boundary of both the erosion $(\bigcup \mathcal{S})^{\ominus \varepsilon}$ and the union $\bigcup \mathcal{B}$. It returns **true** iff there exists an open neighborhood N_v of v such that $N_v \cap (\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq (\bigcup \mathcal{B})$.

Each of those four functions can be implemented in polynomial time and we leave to the reader the choice of an implementation.

```

Input:  $(\mathcal{S}, \varepsilon, \mathcal{B})$ .
if not INCLUSION(BOUNDARY( $\mathcal{S}, \varepsilon$ ),  $\mathcal{B}$ ) then return false;
forall the vertices  $v$  of the boundary of  $\bigcup \mathcal{B}$  do
  | switch POINTLOCATION( $v, \mathcal{S}, \varepsilon$ ) do
  | | case outer return false;
  | | case boundary
  | | | if not LOCALINCLUSION( $v, \mathcal{S}, \varepsilon, \mathcal{B}$ ) then return false ;
  | | | end
  | | endsw
  | end
end
return true ;

```

Algorithm 1: An alternative way to test the inclusion $(\bigcup \mathcal{S})^{\ominus \varepsilon} \subseteq \bigcup \mathcal{B}$.

B Interior of a cone

Lemma 8 *Let $u \in \mathbb{R}^n$ and $C = \{c_1, \dots, c_k\} \subseteq \mathbb{R}^n$. We have the inclusion*

$$\mathring{\text{coni}}(u, C) \subseteq \left\{ u + \sum_{i=1}^k \lambda_i (c_i - u) \mid \lambda_i \in \mathbb{R}, \lambda_i > 0 \right\}$$

In order to prove Lemma 8 we rely on the following technical result:

Lemma 9 *Let $u \in \mathbb{R}^n$ and $C = \{c_1, \dots, c_k\} \subseteq \mathbb{R}^n$. Let $c \in \mathring{\text{coni}}(u, C)$ and $l \in \{1, \dots, k\}$. There exist $\lambda_1, \dots, \lambda_k \geq 0$, $\lambda_l > 0$, such that $c - u = \sum_{i=1}^k \lambda_i (c_i - u)$.*

Proof Let $e_l = (c_l - u)/\|c_l - u\|$. Let $r > 0$ such that $b(c, r) \subseteq \mathring{\text{coni}}(u, C)$. Then $c - r e_l \in \mathring{\text{coni}}(u, C)$ and there are $\mu_1, \dots, \mu_k \geq 0$ such that $c - r e_l - u = \sum_{i=1}^k \mu_i (c_i - u)$. The last expression can be rewritten as follows:

$$c - u = \frac{r}{\|c_l - u\|} (c_l - u) + \sum_{i=1}^k \mu_i (c_i - u)$$

Setting $\lambda_i = \mu_i$ for $i \neq l$, and $\lambda_l = \mu_l + r/\|c_l - u\|$, we obtain that $c - u = \sum_{i=1}^k \lambda_i (c_i - u)$ with $\lambda_i \geq 0$, and $\lambda_l > 0$. \square

We can now prove Lemma 8.

Proof of Lemma 8 By Lemma 9, for each $l \in \{1, \dots, k\}$ there are $\lambda_{l1}, \dots, \lambda_{lk} \geq 0$, $\lambda_{ll} > 0$, such that $c - u = \sum_{i=1}^k \lambda_{li} (c_i - u)$. By summing over $1 \leq l \leq k$ we have

$$k(c - u) = \sum_{i=1}^k \left(\sum_{l=1}^k \lambda_{li} \right) (c_i - u)$$

with $\sum_{l=1}^k \lambda_{li}/k > 0$ for all i . \square

C Gadget for a degree-three vertex when $\delta = 0$

When $\delta = 0$, the set of disks \mathcal{S} computed in Section 4.2 from graph $G = (V, E)$ may fail to satisfy Property (i). In other words, the ε -erosion of the union $S = \bigcup \mathcal{S}$ may not coincide with the rotulae (*i.e.* the centers of disks in \mathcal{S}). To see this, assume G has a degree-three vertex v . The vertex gadget for v consists of three disks whose centers c_1 , c_2 and c_3 form an equilateral triangle. By construction, the circumdisk of triangle $c_1 c_2 c_3$ is centered at v and when $\delta = 0$, the circumdisk has radius ε and is contained in S . Thus, v belongs to the ε -erosion of S without being a rotula and Property (i) does not hold.

In this appendix, we propose a vertex gadget for degree-three vertices, specific to the case $\delta = 0$, while keeping all other gadgets as presented in Section 4.2. More specifically, given a degree-three vertex v adjacent to edges e_1 , e_2 , and e_3 , the new vertex gadget \mathcal{G}_v for v is composed of six disks as depicted in Figure 12: three of them have radius ε and are centered on the three extreme rotulae c_{ve_i} , one for each edge e_i incident to v . These extreme rotulae lie on the boundary of a larger disk of radius 2ε , denoted by \widehat{b}_v . Moreover, a disk slightly smaller than \widehat{b}_v sits just off-center of \widehat{b}_v . We denote it by \overleftarrow{b}_v . Last, a disk of radius ε is centered on point c_v , called vertex rotula, which lies on the boundary of \overleftarrow{b}_v . Notice that

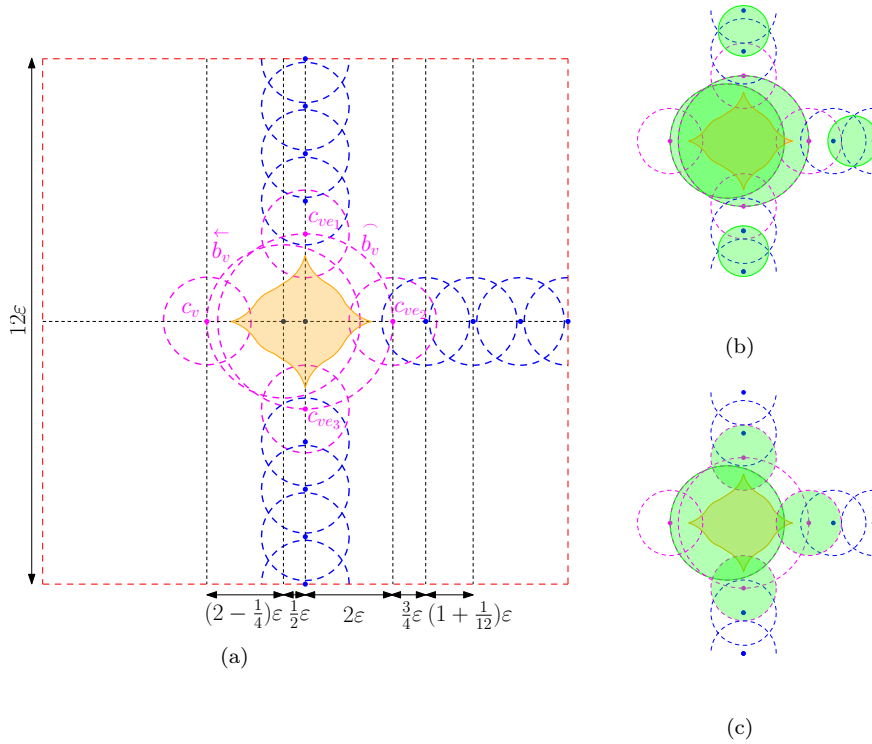


Fig. 12: (a) Gadget for a degree-three vertex when $\delta = 0$; (b)-(c) Illustration of Properties (ix) and (x): the ε -erosion of the vertex gadget is either contained in the union of the two large balls, (b) or in the union of one large ball and the three balls of radius ε centered on extreme rotulae (c).

the vertex gadget for v fits in a square of side length 12ε . (Actually, it fits in a square of side length 8ε but a square of size 12ε allows us to combine it with blocks of Section 4.2.)

Observe that the ε -erosion of the vertex gadget for v is now composed of 5 connected components: the three extreme rotulae c_{ve_i} , the vertex rotula c_v , and an area depicted in orange in Figure 12. The ε -erosion of the vertex gadget for v minus the three extreme rotulae linked to v is referred to as the *core* of v . The core consists of the orange-coloured area together with the vertex rotula c_v . In the following, we call the intersection of the disk $b(c_{ve}, \varepsilon)$ with the core of v the *tip* of the vertex v associated to the edge e and denote it as t_{ve} . We say that a tip t_{ve} and a vertex x are linked to one another whenever $v = x$. Finally, the remainder of v is the core of v minus its three tips. In particular, the remainder contains the vertex rotula c_v .

Given a degree-three vertex v , the key property of its gadget \mathcal{G}_v is that for any collection of disks \mathcal{X} such that $\bigcup \mathcal{X} \subseteq \bigcup \mathcal{G}_v$, if we want to complete \mathcal{X} into a collection of disks that covers the ε -erosion of $\bigcup \mathcal{G}_v$, then two disks are necessary if one of the extreme rotulae or one of the tips linked to v are not covered by disks in \mathcal{X} but one disk is sufficient whenever the three tips and extreme rotulae linked to v are covered by disks in \mathcal{X} ; see Figure 12.

Given as input a graph G and a rational number $\varepsilon \geq 0$, we are now ready to convert G into a collection of rational disks \mathcal{S} so that G admits a vertex cover of cardinality k if and only if $\bigcup \mathcal{S}$ admits a $(0, \varepsilon)$ -ball approximation of cardinality $k + L + D_3$, where D_3 is the number of degree-3 vertices of G . For short, we shall say that a collection \mathcal{S} which

satisfies this property is *valid*. We get a valid collection of disks \mathcal{S} by first drawing G on an orthogonal grid of size 24ε (instead of 8ε as before). We then cover the drawing of G with two types of blocks, each containing at most one vertex gadget or a part of an edge gadget: blocks of size 12ε are used to cover degree-3 vertices and blocks of size 4ε are used to cover edges, degree-0, degree-1 and degree-2 vertices as illustrated in Figure 13.

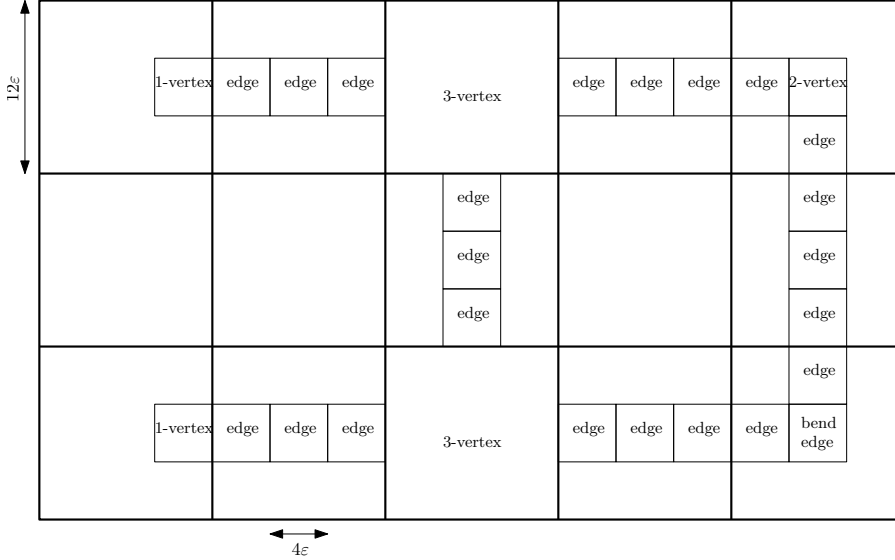


Fig. 13: Conversion of the drawing of graph G in Figure 6.a using blocks of size 12ε and blocks of size 4ε .

We prove below that \mathcal{S} is valid by establishing Properties 9 and 10 which are the equivalent of Properties 4 and 5 in the case $\delta = 0$.

First, let us revisit properties (i), (vi), and (vii) of Section 4.1:

- (i') The ε -erosion of $S = \bigcup \mathcal{S}$ is exactly the collection of rotulae, tips, and remainders.
- (vi') A disk $b \subseteq S^{\oplus \delta}$ intersects the ε -erosion of at most one vertex gadget v .
- (vii') Let v be a vertex. There exists a disk $b \subseteq S$ that covers all extreme rotulae and all tips linked to v .

Two extra properties are needed:

- (ix) Let v be a degree-three vertex. There exists a disk $b \subseteq S$ that covers the remainder of v .
- (x) Let v be a degree-three vertex. For any edge e incident to v , (a) there exists a disk b that covers the extreme rotula c_{ve} , the tip t_{ve} , and the regular rotulae neighbour to c_{ve} ; (b) there exists a disk b that covers the extreme rotula c_{ve} , the tip t_{ve} , and no regular rotulae.

Note that an example of disk b that satisfies Property (vii') is \widehat{b}_v , an example of disk b that satisfies Property (ix) is \overleftarrow{b}_v (see Figure 12(c)). Moreover, the three disks $b(c_{ve}, \varepsilon)$ (small green disks on Figure 12(c)) satisfy Property (x)(a).

Recall that for each edge e incident to a vertex v , we introduced a canonical covering that covers c_{ve} and all the regular rotulae of edge e . Whenever v is a degree-three vertex, we furthermore enforce this canonical covering to contain the ball $b(c_{ve}, \varepsilon)$ (see Figure 14).

We are now ready to rewrite Properties 4 and 5 as follows:




Fig. 14: Canonical covering for an extreme rotula linked to a degree-three vertex.

Property 9 If G has a vertex cover $C \subseteq V$, then S has a $(0, \varepsilon)$ -ball approximation \mathcal{B} with $|\mathcal{B}| = L + D_3 + |C|$, where D_3 is the number of degree-three vertices of G .

Proof The principle is the same as in the proof of Property 4: for each vertex $v \in C$ we start by selecting disks that cover the ε -erosion of the vertex gadget and the extreme rotulae linked to v . If we denote by $D_3(C)$ the number of degree-three vertices in C , by properties (vii') and (ix), this gives a set \mathcal{C} of $|C| + D_3(C)$ disks that cover at least one extreme rotula per edge.

Then, as before, all the regular rotulae can be covered with L disks using the appropriate canonical covering for each edge. Consider now a degree-three vertex $v \notin C$. For any edge e incident to v , the canonical covering of e covers the extreme rotula c_{ve} and the tip linked to c_{ve} . By Property (ix), \mathcal{C} can be completed into a $(0, \varepsilon)$ -ball approximation with one disk for each degree-three vertex $v \notin C$, thus $D_3 - D_3(C)$ disks (see Figure 12(b)), which concludes the proof.

Property 10 If S has a $(0, \varepsilon)$ -ball approximation \mathcal{B} , then G has a vertex cover C with $|C| \leq |\mathcal{B}| - L - D_3$, where D_3 is the number of degree-three vertices of G .

Proof Similarly to the proof of Property 5, we proceed to substitutions to transform any $(0, \varepsilon)$ -ball approximation \mathcal{B} into another one called \mathcal{R} having the property that it contains a canonical covering of each edge, and exactly one disk for each vertex rotula.

As before, we denote by \mathcal{C}_{ve} the canonical covering of edge e that covers the extreme rotulae c_{ve} . Initializing \mathcal{R} to \mathcal{B} , we modify \mathcal{R} in two steps as follows:

1. For each degree-three vertex v , consider the set of disks \mathcal{B}_v of \mathcal{B} that cover c_v . Replace \mathcal{B}_v by the singleton disk $\{b_v\}$. Recall that b_v satisfies Property (ix), and covers the remainder of v .
2. For each edge $e = \{u, v\}$, consider the set of disks \mathcal{B}_e of \mathcal{B} that contain at least one regular rotula of e . We build another set of disks \mathcal{R}_e that covers all regular rotulae of e , with $|\mathcal{R}_e| = \ell(e)$ and proceed to the following substitutions:
 - if \mathcal{B}_e covers one extreme rotula or none, \mathcal{R}_e is set to the appropriate canonical covering. We replace \mathcal{B}_e by \mathcal{R}_e .
 - if \mathcal{B}_e covers the two extreme rotulae, then $|\mathcal{B}_e| \geq \ell(e) + 1$. Choose arbitrarily one of the two extreme rotulae, say c_{ve} . If the other vertex u is not of degree three, choose $b \subseteq S$ that covers c_{ue} but no regular rotula. If u is a degree-three vertex, remember that, by definition, any disk b' of \mathcal{B}_e contains a regular rotula. Thus, if b' contains also c_{ue} , the part of S that b' covers is contained in the union of the tip t_{ue} with the remainder of u . The remainder of u being covered by b_u thanks to step 1, we can replace any such disk b' by a disk $b \subseteq S$ that covers c_{ue} , t_{ue} but no regular rotula and still preserve the $(0, \varepsilon)$ -ball approximation property. This disk exists thanks to Property (x)(b). We replace \mathcal{B}_e by $\mathcal{C}_{ve} \cup \{b\}$, and \mathcal{R}_e is equal to \mathcal{C}_{ve} .

Each of these substitutions preserves the $(0, \varepsilon)$ -ball approximation property and does not increase the cardinality of the resulting collection of disks. Consider the set \mathcal{C} of disks of \mathcal{R} that do not contain any regular rotula, nor any vertex rotula. We have:

$$\mathcal{C} \subseteq \mathcal{R} \setminus \left(\bigcup_{e \in E} \mathcal{R}_e \cup \bigcup b_v \right),$$

where v ranges over the set of degree-three vertices. By construction, $|\mathcal{C}| \leq |\mathcal{R}| - L - D_3 \leq |\mathcal{R}| - L - D_3$ and \mathcal{C} covers at least one extreme rotula per edge. The end of the proof is similar to the proof of Property 5. Let C be the set of vertices $v \in V$ linked to an extreme rotula covered by a disk in \mathcal{C} . By construction, any $b \in \mathcal{C}$ intersects the ε -erosion of exactly one vertex gadget (by Property (iv')). It is thus possible to map each $b \in \mathcal{C}$ to one vertex $v \in C$, showing that $|C| \leq |\mathcal{C}|$. \square

To conclude, note that all the disks of this construction have rational centers and radii. Moreover, this construction can actually be extended for any value $\delta \geq 0$.² Being rather more complex, we chose to present it in Appendix, and keep a simpler construction in the article.

² See the Geogebra file http://www.gipsa-lab.fr/~isabelle.sivignon/Code/degree-three_vertex_gadget_forall_delta.ggb that provides a parameterized construction for any value of δ and ε with $\delta + \varepsilon > 0$.