ε-covering is NP-complete

Dominique Attali* Tuong-Bach Nguyen† Isabelle Sivignon‡

Abstract

Consider the dilation and erosion of a shape S by a ball of radius ε. We call ε-covering of S any collection of balls whose union lies between the dilation and erosion of S. We prove that finding an ε-covering of minimum cardinality is NP-complete, using a reduction from vertex cover.

1 Introduction

Unions of balls are common shape representations, useful for instance to describe molecules in biochemistry [3], to quickly detect collisions [1] between shapes or to derive higher-level representations such as medial axis descriptions. The ubiquity of unions of balls is largely due to the existence of provably good conversion algorithms that allow us to derive them from various representations such as point clouds and polygonal meshes [7]. In that case, the union of balls output by the conversion process generally provides only an approximation of the original shape. The quality of the approximation can be measured by its geometric error. In this paper, we introduce a parameter ε that controls the admissible geometric error in a novel way. We say that a collection of balls provides an ε-covering of a shape if its union is contained in the dilation of the shape by b_ε and contains the erosion of the shape by b_ε, where b_ε refers to the ball with radius ε centered at the origin. We are interested in the problem of computing such a collection of balls with minimum cardinality. Our main result is that this problem is NP-complete.

2 Statement of result

In this paper, we suppose that $\mathbb{R}^d$ is endowed with the Euclidean distance. For any point c and real $r > 0$, we denote by $b(c, r)$ the open ball centered at c with radius r. Let $S \subseteq \mathbb{R}^d$ and $\varepsilon > 0$ a real number. The **dilation** of S (by ε) is $S^{\oplus\varepsilon} = \cup_{x\in S} b(x, \varepsilon)$. The **erosion** of S (by ε) is $S^{\ominus\varepsilon} = \{x \mid b(x, \varepsilon) \subseteq S\}$. For any collection of balls $\mathscr{B}$, we adopt the notation $\bigcup\mathscr{B} = \cup_{b\in\mathscr{B}} b$. A collection of balls is rational if each of its balls has a rational radius and a center with rational coordinates. We also assume that these rationals can be represented by integers bounded by a given constant.

Definition 1 *An ε-covering of S is a collection of balls $\mathscr{B}$ such that $S^{\ominus\varepsilon} \subseteq \bigcup\mathscr{B} \subseteq S^{\oplus\varepsilon}$. Additionally if $\mathscr{B}$ is rational, it is a rational ε-covering.*

We are interested in ε-coverings that achieve minimum cardinality and must solve this problem:

Problem 1 (Rational ε-covering problem) *Let $\mathscr{S}$ be a finite rational collection of balls in $\mathbb{R}^d$, $\varepsilon > 0$ a rational and k a positive integer. Does $\bigcup\mathscr{S}$ have a rational ε-covering with at most k balls?*

The purpose of this paper is to prove the following:

Theorem 1 *The rational ε-covering problem in $\mathbb{R}^2$ is NP-complete.*

To prove our theorem, we need to establish that the ε-covering problem is both in NP and NP-hard. For lack of space, we only sketch the proof of the NP property in Section 3 and focus on the NP-hardness in Section 4. Indeed, a formal proof of the former property requires several technical results we will not detail. Note however that the purpose of the rational and $\mathbb{R}^2$ restrictions is to clear these technical hurdles. As for the latter property, it is obtained through a reduction from a variant of the vertex cover problem, which remains valid in all dimensions.

3 NP

The ε-covering problem is in NP if one can verify in polynomial time in the size of $\mathscr{S}$ whether a collection of balls is a solution. Given a shape $S = \bigcup\mathscr{S}$, a parameter $\varepsilon > 0$ and some collection of balls $\mathscr{B}$, is it possible to check the two inclusions $S^{\ominus\varepsilon} \subseteq \bigcup\mathscr{B} \subseteq S^{\oplus\varepsilon}$ in polynomial time? Since we only address the $\mathbb{R}^2$ case in this section, balls are simply disks whose boundaries are circles. We propose a method that relies on the arrangement of those boundary circles.

In general, the exact computation of these arrangements requires an exact handling of real numbers. Though, with our restriction to rationals, we only need to handle algebraic numbers. This can be done using the *isolating interval representation* [8].

An arrangement of circles is the subdivision of $\mathbb{R}^2$ into open connected cells which is induced by these circles. If we consider the original disk supported by

*Gipsa-lab, dominique.attali@gipsa-lab.grenoble-inp.fr
†Gipsa-lab, tuong-bach.nguyen@gipsa-lab.grenoble-inp.fr
‡Gipsa-lab, isabelle.sivignon@gipsa-lab.grenoble-inp.fr

This is an extended abstract of a presentation given at EuroCG 2016. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear in a conference with formal proceedings and/or in a journal.

each circle, each cell can then be characterized by two sets of disks: the disks that contain the cell, and those that do not. Along with the arrangement itself, this information can be computed in polynomial time for each cell [5, 4]. Owing to this, one can test whether a particular disk b is contained in a finite union of disks $\bigcup \mathscr{S}$. Indeed, it suffices to compute the arrangement of $\mathscr{S} \cup \{b\}$ and then check for every cell covered by b if it is also covered by some disk of $\mathscr{S}$. This can obviously be extended to verifying if a finite union of disks contains another such union. Since $S^{\oplus\varepsilon}$ is the union of the dilated balls of $\mathscr{S}$, we can test if $\mathscr{B}$ satisfies $\bigcup \mathscr{B} \subseteq S^{\oplus\varepsilon}$. As for the second condition, it requires the following technical result.

Proposition 2 *Let $\mathscr{V}$ be the set of vertices of the boundary of S. Consider $\mathscr{S}' = \{b^{\ominus\varepsilon} \mid b \in \mathscr{S}\} \cup \{b(v,\varepsilon) \mid v \in \mathscr{V}\}$. Then there exists a collection $\mathscr{C}$ of cells of the arrangement of $\mathscr{S}'$ such that $S^{\ominus\varepsilon} = \bigcup \mathscr{C}$, and these cells can be found in polynomial time.*

Thus, by computing the arrangement of $\mathscr{S}' \cup \mathscr{B}$ we can test whether $S^{\ominus\varepsilon} \subseteq \bigcup \mathscr{B}$ by inspecting the cells that belong to $S^{\ominus\varepsilon}$. Though a generalization of Proposition 2 to higher dimension seems feasible, the computation of the arrangement of $\mathscr{S}' \cup \mathscr{B}$ is not straightforward.

4 NP-hardness

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 $F \subseteq V$ is a vertex cover of G if every edge of E is incident to a vertex of F. Finding a minimum vertex cover is NP-hard, even when restricted to cubic planar graphs [6]. We shall perform the reduction from this particular variant. Let $G = (V, E)$ be a planar graph of degree at most 3. For any $\varepsilon > 0$, we show how to build a finite collection of balls $\mathscr{S}(G, \varepsilon)$ such that G has a vertex cover of cardinality k if and only if $S(G, \varepsilon) = \bigcup \mathscr{S}(G, \varepsilon)$ has an ε-covering of cardinality $k + N$, where N is a constant depending on $\mathscr{S}(G, \varepsilon)$. To simplify notations, we shall refer to $\mathscr{S}(G, \varepsilon)$ and $S(G, \varepsilon)$ simply as $\mathscr{S}$ and S.

4.1 Reduction from vertex cover

Our construction of $\mathscr{S}$ uses two types of balls: rotula balls and ghost balls. **Rotula** balls are balls of radius ε; their centers are called rotulae. They are used in finite sequences of odd length at least 3 that we will call edge gadgets. Two subsequent rotulae of an edge gadget are said to be neighbours of each other. We distinguish two types of rotulae: linking rotulae that are the endpoints of each edge gadget and have only one neighbour, and regular rotulae which are all other rotulae and have two neighbours. See Figure 1b. Besides rotula balls, our construction uses a second type of balls, the **ghost** balls. These balls have radii $\lambda\varepsilon$, $\lambda < 1$; their centers are called ghosts. A single one of these ghost balls constitutes a vertex gadget, which is connected to different edge gadgets through one of their linking rotulae. Thus each linking rotula is associated with one unique ghost, whereas a ghost may be associated with up to 3 linking rotulae depending on the degree of the vertex it was converted from.

Thus, each edge (resp. vertex) in G is converted into an edge gadget (resp. vertex gadget) and we define $\mathscr{S}$ as the collection of all rotula balls and ghost balls resulting from that conversion. At this point, we haven't yet specified the number of rotula balls per edge that we need (only that it should be an odd number) nor the location of rotulae and ghosts. This will be done in Section 4.2 where we build $\mathscr{S}$ from an orthogonal grid drawing of G so that it fulfills the properties below; see Figure 1 for an example.

(i) The erosion of S is exactly the collection of rotulae, $S^{\ominus\varepsilon} = \{c \mid c \text{ is a rotula}\}$.

(ii) Any ball $b \subseteq S^{\oplus\varepsilon}$ contains at most 2 regular rotulae.

(iii) If a ball $b \subseteq S^{\oplus\varepsilon}$ contains 2 regular rotulae, then they are neighbours and b does not contain any other rotula, neither regular nor linking.

(iv) Let c be a regular rotula, c_+ and c_- its two neighbours. There exist two balls $b_+, b_- \subseteq S^{\oplus\varepsilon}$ such that $\{c_+, c\} \subseteq b_+$ and $\{c_-, c\} \subseteq b_-$.

(v) Any ball $b \subseteq S^{\oplus\varepsilon}$ contains at most 3 linking rotulae.

(vi) If a ball $b \subseteq S^{\oplus\varepsilon}$ contains 2 or 3 linking rotulae, then these linking rotulae are associated with the same ghost and b only contains linking rotulae associated with this ghost.

(vii) Let c_g be a ghost. There exists a ball $b \subseteq S^{\oplus\varepsilon}$ that contains all linking rotulae associated with c_g.

(viii) If a ball $b \subseteq S^{\oplus\varepsilon}$ contains both a regular rotula and a linking rotula, then these are neighbours and b does not contain any other rotula, neither regular nor linking.

Henceforth, we shall make no distinction between an edge of G and its conversion into an edge gadget, and likewise for a vertex of G and its corresponding ghost. Recall that an edge gadget $e \in E$ is a finite sequence of rotula balls of odd length at least 3. We denote by $n(e) \geq 1$ the integer such that $2n(e) + 1$ is the length of the edge gadget e. Thus, e has $2n(e) - 1$ regular rotulae. From property (ii) we need at least $\lceil (2n(e) - 1)/2 \rceil = n(e)$ balls in $S^{\oplus\varepsilon}$ in order to cover these regular rotulae. By (iv) there always exists a

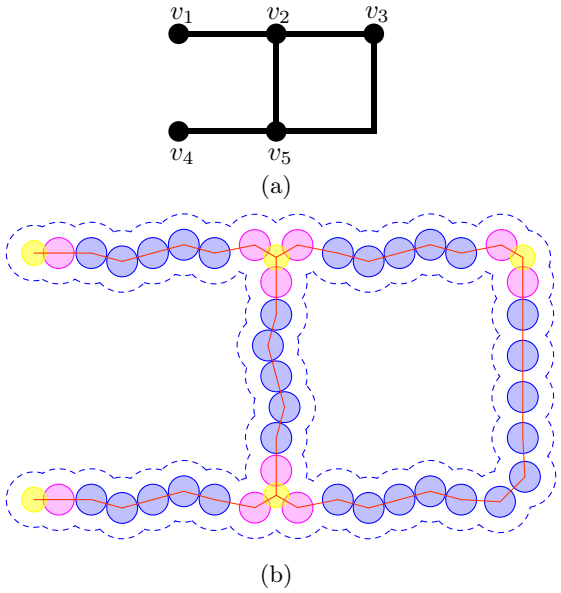

Figure 1: An example of conversion from (a) an orthogonal grid drawing of a graph with 5 vertices and 5 edges to (b) a good collection of balls: blue disks are regular rotulae balls, purple are linking rotulae balls and yellow are ghosts balls. The dilation is bounded by the dashed blue lines. All figures use ghost balls with radius $\lambda\varepsilon$, $\lambda = 0.8$.

collection of $n(e)$ balls in $S^{\oplus\varepsilon}$ which covers these rotulae plus one of the two linking rotulae of e, and that linking rotula can be chosen arbitrarily. Indeed, it suffices to cover pairs of neighbouring rotulae of e with non-overlapping balls in $S^{\oplus\varepsilon}$, while making sure that the linking rotula to cover and its neighbour are one of these pairs (see Figure 2). This gives two possible coverings of regular rotulae of e (one for each linking rotulae) which we shall refer to as **canonical** for e. Furthermore, properties (iii) and (viii) guarantee that any ball containing a regular rotula will only contain rotulae belonging to the same edge gadget. Therefore it is necessary and sufficient to use $n(e)$ balls to cover the regular rotulae of an edge gadget e, and these $n(e)$ balls exclusively cover rotulae of e. However, to entirely cover an edge gadget, one extra ball is required to cover the second linking rotula, for a total of $n(e)+1$ balls. Contrary to the previous $n(e)$ balls, by (v) that extra ball may be shared among several edge gadgets to cover their last linking rotulae. We define $N = \sum_{e\in E} n(e)$. N is the number of balls needed to cover all regular rotulae with canonical coverings.

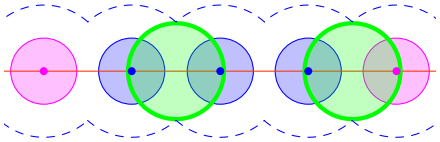

Figure 2: A canonical covering of an edge gadget (green disks). Same color convention as in Figure 1b.

Proposition 3 *If G has a vertex cover $F \subseteq V$, then S has an ε-covering $\mathscr{B}$ with $|\mathscr{B}| = N + |F|$.*

Proof. For each vertex $u \in F$, we use property (vii) and select a ball covering all linking rotulae associated with u. By the vertex cover property, this yields $|F|$ balls that cover at least one linking rotula per edge. Using the appropriate canonical coverings of each edge, we then complete the ε-covering with N more balls to cover the regular rotulae and any remaining linking rotulae. By (i), this collection of balls is an ε-covering since it contains every rotula. □

Proposition 4 *If S has an ε-covering $\mathscr{B}'$, then G has a vertex cover F with $|F| \leq |\mathscr{B}'| - N$.*

Proof. Without loss of generality, we may assume that all balls in $\mathscr{B}'$ cover at least one rotula. Indeed, if a ball b does not cover any rotula, it can be removed from $\mathscr{B}'$ while keeping the property that $\mathscr{B}'$ is an ε-covering. Starting from $\mathscr{B}'$, we first deduce an ε-covering $\mathscr{B}$ 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 $\mathscr{H}(e, \mathscr{B}') = \{b \in \mathscr{B}' \mid \text{b contains a regular rotula of } e\}$. Note that for different edges the $\mathscr{H}(e, \mathscr{B}')$ are disjoint and that $|\mathscr{H}(e, \mathscr{B}')| \geq n(e)$. Given a linking rotula u of $e \in E$, we denote by $\mathscr{C}_e(u)$ the canonical covering that contains u and all the regular rotulae of e. Initializing $\mathscr{B}$ to $\mathscr{B}'$, we then transform $\mathscr{B}$ as follows, replacing each $\mathscr{H}(e, \mathscr{B}')$ according to the following three cases:

- $\bigcup \mathscr{H}(e, \mathscr{B}')$ contains 0 linking rotula. We choose an arbitrary linking rotula u and replace $\mathscr{H}(e, \mathscr{B}')$ with $\mathscr{C}_e(u)$.
- $\bigcup \mathscr{H}(e, \mathscr{B}')$ contains 1 linking rotula u. Then we simply replace $\mathscr{H}(e, \mathscr{B}')$ with $\mathscr{C}_e(u)$.
- $\bigcup \mathscr{H}(e, \mathscr{B}')$ contains both linking rotulae. Then, $|\mathscr{H}(e, \mathscr{B}')| \geq n(e) + 1$. We choose an arbitrary linking rotula u, and let b be a ball containing the other linking rotula but no regular rotula. We replace $\mathscr{H}(e, \mathscr{B}')$ with $\mathscr{C}_e(u) \cup \{b\}$.

Each of these substitutions preserves the ε-covering property and does not increase the cardinality of the resulting collection of balls. Consider the balls of $\mathscr{B}$ that do not contain any regular rotula, $\mathscr{F} = \mathscr{B} \setminus (\cup_{e\in E} \mathscr{H}(e, \mathscr{B}))$. By construction, $|\mathscr{F}| = |\mathscr{B}| - N \leq |\mathscr{B}'| - N$. Let $F = \{u \in V \mid \exists b \in \mathscr{F},\ b \text{ contains a linking rotula associated with the ghost } u\}$. We claim that F is a vertex cover of G and that its cardinality satisfies $|F| \leq |\mathscr{F}|$. All $b \in \mathscr{F}$ must contain at least one linking rotula, thus F is empty if and only if $\mathscr{F}$ is empty. In this particular case, the empty set is a vertex cover of G: indeed, G must have no edges because otherwise $\mathscr{B}$ would only cover half of the linking rotulae. Assume now that $\mathscr{F}$ is not empty. By (vi),

any $b \in \mathscr{F}$ yields at most one vertex in F. As for the vertex cover property, recall that $\cup_{e \in E} \mathscr{H}(e, \mathscr{B})$ covers exactly all regular rotulae and one linking rotula per edge. Hence $\mathscr{F}$ must cover the remaining linking rotula of each edge. The definition of F thus ensures it contains at least one endpoint of each edge. □

4.2 Practical construction of $\mathscr{S}$

All that remains is to build a collection $\mathscr{S}$ fulfilling properties (i) to (viii) given G and ε. To do so, we rely on the following result.

Theorem 5 ([2]) *There is a linear time and space algorithm to draw a connected at most cubic graph on an orthogonal grid.*

Given such a drawing of G, we now describe a way to convert it into an appropriate collection of balls. We rely on the orthogonal drawing in Figure 1a as an example. To perform the conversion, we fix the size of the grid to 16ε so that we can fit square blocks of size $8\varepsilon \times 8\varepsilon$ as in Figure 3. There are two different ways in

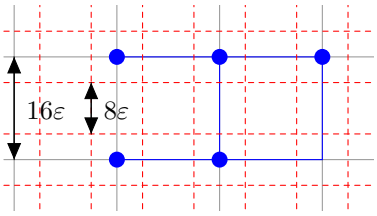

Figure 3: Grid division into blocks. Gray lines represent the grid, dashed red lines are the blocks and the example graph is in blue.

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 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 balls that satisfies properties (i) to (viii). For blocks containing a vertex, see Figure 4 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 5.

5 Conclusion

Finding an ε-covering of minimum cardinality is NP-complete. Though the proof presented here relies on a family of shapes with many connected components and whose erosion does not preserve the genus, the

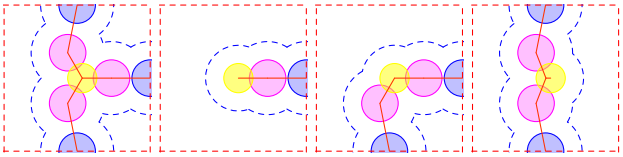

Figure 4: Block conversions for a vertex of degree (a) 3, (b) 1, (c) 2 in a bend and (d) 2 in a straight line. The red dashed square delimit the block. Same color convention as in Figure 1b.

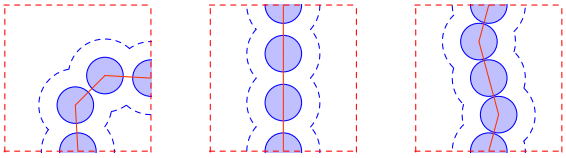

Figure 5: Block conversions for (a) a bent edge, (b) an even and (c) odd straight edge. Same color convention as in Figure 4.

result does not depend on it. Indeed, a similar albeit more intricate construction shows that the problem is still hard when restricted to connected shapes that are homotopy equivalent to their erosion.

References

[1] G. Bradshaw and C. O'Sullivan. Adaptive medial-axis approximation for sphere-tree construction. *ACM Transactions on Graphics (TOG)*, 23(1):1–26, 2004.

[2] T. Calamoneri and R. Petreschi. An efficient orthogonal grid drawing algorithm for cubic graphs. In *Comp. and Combinatorics*, pages 31–40. Springer, 1995.

[3] F. Cazals, T. Dreyfus, S. Sachdeva, and N. Shah. Greedy geometric algorithms for collection of balls, with applications to geometric approximation and molecular coarse-graining. In *Computer Graphics Forum*, volume 33, pages 1–17. Wiley Online Lib., 2014.

[4] F. Cazals and S. Loriot. Computing the arrangement of circles on a sphere, with applications in structural biology. *Computational Geometry*, 42(6):551–565, 2009.

[5] H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir. Arrangements of curves in the plane—topology, combinatorics, and algorithms. *Theoretical Computer Science*, 92(2):319–336, 1992.

[6] M. R. Garey and D. S. Johnson. The rectilinear steiner tree problem is NP-complete. *SIAM Journal on Applied Mathematics*, 32(4):826–834, 1977.

[7] B. Miklos, J. Giesen, and M. Pauly. Discrete scale axis representations for 3d geometry. In *ACM Transactions on Graphics (TOG)*, volume 29, page 101. ACM, 2010.

[8] C.-K. Yap. Towards exact geometric computation. *Computational Geometry*, 7(1):3–23, 1997.