Finding a Minimum Medial Axis of a Discrete Shape is NP-hard

David Coeurjolly [a], Jérôme Hulin [b,1], Isabelle Sivignon [a]

[a] *LIRIS, UMR CNRS 5205, Université Claude Bernard Lyon 1, F-69622 Villeurbanne, France*

[b] *LIF, UMR CNRS 6166, Université de la Méditerranée, 163, av. de Luminy F-13288 Marseille, France*

Abstract

The medial axis is a classical representation of digital objects widely used in many applications. However, such a set of balls may not be optimal: subsets of the medial axis may exist without changing the reversivility of the input shape representation. In this article, we first prove that finding a minimum medial axis is a NP-hard problem for the Euclidean distance. Then, we compare two algorithms which compute an approximation of the minimum medial axis.

Key words: Minimum Medial Axis, NP-completeness, bounded approximation algorithm.

1 Introduction

In binary images, the *Medial Axis* (MA) of a shape $\mathcal{S}$ is a classic tool for shape analysis. It was first proposed by Blum [2] in the continuous plane; then it was defined by Pfaltz and Rosenfeld in [14] to be the set of centers of all maximal disks in $\mathcal{S}$, a disk being maximal in $\mathcal{S}$ if it is not included in any other disk in $\mathcal{S}$. This definition allows the medial axis to be computed in a discrete framework, i.e., if the working space is the rectilinear grid $\mathbb{Z}^n$. The medial axis has the property of being a *reversible* coding: the union of the disks of MA($\mathcal{S}$) is exactly $\mathcal{S}$.

In order to compute the medial axis of a given discrete shape $\mathcal{S}$, we first proceed by computing the *Distance Transform* (DT) of $\mathcal{S}$. The distance trans-

[1] Supported in part by ANR grant BLAN06-1-138894 (projet OPTICOMB)

form is a bitmap image in which each point is labelled with the distance to the closest background point. For either d_4, d_8, any given chamfer distance or the Euclidean distance d_E, the distance transform can be computed in linear time with respect to the number of grid points [18,4,7,11]. For the simple distances d_4 and d_8, MA is extracted from DT by picking the local maxima in DT [18,4,16].

Polynomial time algorithms exist to extract MA from DT in the case of the chamfer norms or the Euclidean distance [16,17]. A Reduced Medial Axis (RMA) is presented in [8]: it is a reversible subset of the medial axis, that can be computed in linear time. Despite the fact that the medial axis exactly describes the shape $\mathcal{S}$, it may not be a set with minimum cardinality of balls covering $\mathcal{S}$: indeed, a maximal disk of the medial axis covered by a union of maximals disks is not necessary for the reconstruction of $\mathcal{S}$.

In this article, we investigate the minimum medial axis problem that aims at defining a set of maximal balls with minimum cardinality which cover $\mathcal{S}$. This problem has already been addressed with algorithms that experimentally filter the medial axis [5,6,15,6,13].

In section 2 we first detail some preliminaries and the fundamental definitions used if the rest of the paper. Section 3 presents the proof that the minimum medial axis problem is NP-hard. Finally, we compare the results given a greedy algorithm with the approximation algorithm proposed in [15] (Section 4).

2 Preliminaries and Related Results

First of all, we remind definitions related to the discrete medial axis. Given a metric d, a (open) ball B of radius r and center p is the set of grid points q such that $d(p, q) < r$. In the following, we consider the Euclidean metric and extension of the results to other metric (such as Chamfer masks for example) will be discussed in section 5.

Definition 1 (Maximal ball) *A ball B is* maximal *in a discrete shape $\mathcal{S} \subseteq \mathbb{Z}^n$ if $B \subseteq \mathcal{S}$ and if B is not entirely covered by another ball contained in $\mathcal{S}$.*

Based on this definition, the medial axis is given by:

Definition 2 (Medial axis) *The medial axis of a shape $\mathcal{S} \subseteq \mathbb{Z}^n$ is the set of all maximal balls in $\mathcal{S}$.*

For the rest of the paper, we focus on the dimension 2. By definition, the medial axis of a shape $\mathcal{S}$ is a reversible encoding of $\mathcal{S}$. Indeed given the centers and the radii associated to the medial axis balls, one can reconstruct entirely the input shape $\mathcal{S}$ (this process is called the Reverse Distance Transformation [18,3,4,19,8]).

However, this representation is not minimum in the number of balls as illustrated in Figure 1: the set of balls with highlighted centers in the left shape corresponds to the medial axis given by Definition 2. However, if we consider the subset of the medial axis depicted in the right figure, we still have a reversible description of the shape with less balls.

1	1	1	1	1	1	1
1	④	④	④	④	④	1
1	1	1	1	1	1	1

1	1	1	1	1	1	1
1	④	4	④	4	④	1
1	1	1	1	1	1	1

Fig. 1. (*Left*) highlighted points correspond to the centers of the medial axis balls for the Euclidean metric (squared values), (*right*) a subset of the medial axis whose balls still cover the entire shape.

In this paper, we address the problem of finding a subset of the medial axis that still covers all points of $\mathcal{S}$.

In computational geometry, covering a polygon with a minimum number of a specific shape (*e.g.* convex polygons, squares, rectangles,...) usually leads to NP-complete or NP-hard problems [10]. From the literature, a related result proposed in [1] concerns the minimum decomposition of an orthogonal polygon into squares. At first sight, this result seems to be closely related to the minimum medial axis computation based on the d_8 metric. However, in the discrete case, d_8 balls are centered on grid points and thus have odd widths. Due to this specificity, results of [1] cannot be used neither for the d_8 nor the Euclidean metrics. However, the proof given in the following sections is inspired by this related work.

3 NP-hardness of the Minimum Medial Axis Problem

Definition 3 (k-Medial Axis Problem (k-MA)) *Given a discrete shape $\mathcal{S} \subseteq \mathbb{Z}^2$ of finite cardinality and a positive integer k, is it possible to cover exactly $\mathcal{S}$ with at most k (possibly overlapping) discrete maximal balls ?*

In order to prove the NP-hardness of k-MA, we use a polynomial reduction of the Planar-4 3-SAT problem. Let $\phi(V, C)$ be the boolean formula in Conjunctive Normal Form (CNF) consisting of a list C of clauses over a set V of variables. The *formula-graph* $G(\phi(V, C))$ of a CNF formula $\phi(V, C)$ is the bipartite graph in which each vertex is either a variable $v \in V$ or a clause $c \in C$, and there is an edge between a variable $v \in V$ and a clause $c \in C$ if v occurs in *c.* A *Planar 3-SAT* formula ϕ is a CNF formula for which the formula-graph $G(\phi)$ is planar and each clause is a 3-clause (i.e., a clause having exactly 3 literals).

In the following, we prefer a reduction based on the Planar-4 3-SAT problem: an instance of this problem is an instance of Planar 3-SAT such that the degree of each vertex, in the formula-graph, associated to a variable is bounded by 4. In other words, a variable may appear at most four times in the boolean formula.

Definition 4 (Planar-4 3-SAT Problem) *Given a Planar-4 3-SAT formula $\phi(V, C)$, does there exist a truth assignement of the variables in V which satisfies all the clauses in C ?*

Planar-4 3-SAT was shown to be NP-complete in [12].

The reduction from any given Planar-4 3-SAT formula ϕ to an instance of k-MA consists in constructing a discrete shape $\mathcal{S}(\phi)$ and finding an integer $k(\phi)$ in polynomial time such that ϕ is satisfiable if and only if $\mathcal{S}(\phi)$ can be covered by $k(\phi)$ balls.

3.1 Variables

Let us first consider the geometrical interpretation of variables. Figure 2 presents a 4-connected discrete object, so called *variable gadget* in the following, defined by the set of grid points below the horizontal dashed line. We call the *extremities* of the variable gadget, the eight vertical parts of the gadget of width 3, numbered on Figure 2. These extremities are used to plug the "wires" that represent the edges of a formula-graph.

Any minimum covering of this object has 72 balls. None of these minimum coverings allow protrusions from both one odd extremity and one even extremity. However, one minimum covering allows balls to protrude out at all odd extremities by one row of grid points (Figure 2 top); while another minimum covering allows balls to protrude out at all even extremities also by one row of grid points (Figure 2 bottom). These two coverings mimic the two possible truth assignements of a variable. Without loss of generality, the first covering will correspond to a True assignment, and the other one to a False assignment of the variable.

If the gadget represents the variable x, then each odd extremity carries the literal x, while each even extremity carries the literal $\bar{x}$. A protrusion from a variable extremity can be viewed as a signal 'True' sent from the variable to the clauses. Thus, wires which are used to connect variables and clauses are plugged on odd extremities for positive literals and on even extremities for negative literals.

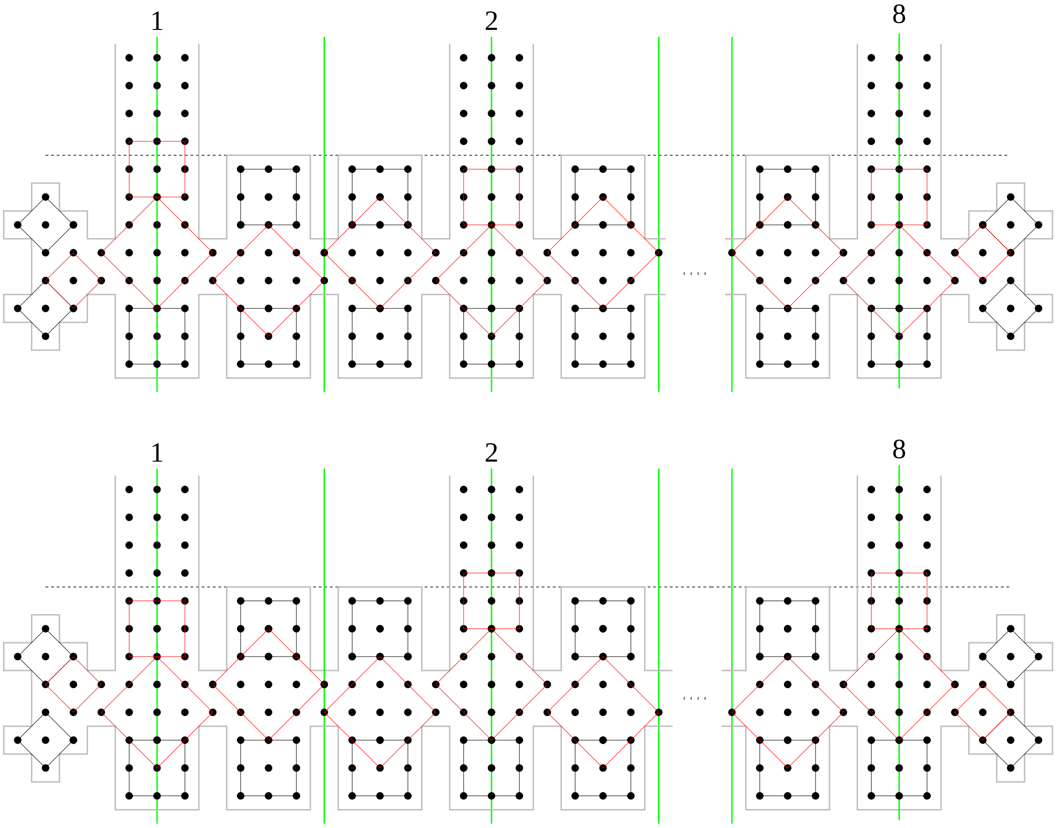

Fig. 2. Two minimum coverings of a variable gadget (from left to right) $x, \bar{x}, \bar{x}$; a True assignement of the variable x (top), and False assignement (bottom).

Note that this object and its decomposition ares invariant under rotation of angle $\frac{\pi}{2}$. Furthermore, the extremities are centered on abscissas with constant values modulo 6 (represented by vertical lines of Figure 2). This modulo operation on the coordinates will be used to align the objects and to connect them to each other.

3.2 Wires

In order to connect variables to clauses, we need wires that correspond to edges in the embedding of the formula-graph. A wire must be designed such that it carries the 'True' signals (protrusions) and 'False' signals (no protrusion) from variable extremities to clauses without altering the signal (see Fig. 3). In this case, we can define straight wires made of $3 \times N$ sets of grid points. If $N \equiv 0 \mod 3$, then the signal sent at the left extremity of the wire will be propagated to the right extremity. Furthermore a wire can be bent at angle $\frac{\pi}{2}$ (see Fig. 3). In this case, two minimum decompositions still exist such that if a ball protrudes from one extremity of the wire, then another ball also protrudes out at the other extremity. Furthermore, straight wires and bends can be designed such that the alignment of the abscissa and ordinates of the shape is preserved (*i.e.* is constant modulo 3).

Now, if we consider a complex wire with straight parts and bends, the signals are propagated during the construction of the minimum covering from on extremity to the other one (proof by induction on the number of bends and straight parts).

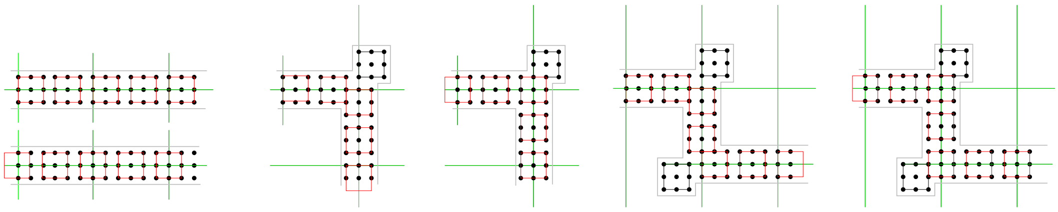

Fig. 3. Wires carrying 'True' or 'False' signals - from left to right: a straight wire, a simple bend, a shift.

3.3 Clauses

Finally, we introduce the *clause gadget*, a component that geometrically simulates a clause. This gadget is the set of grid points at the right of the vertical dashed line in Fig.4. Note that this gadget is not symmetrical because we shall not allow an open ball of radius $\sqrt{8}$ to be placed in its center.

Independently covering this gadget requires at least 10 balls (see Fig.4, left). However, if one open ball of radius 2 is protruding from some wire by one column, carrying a 'True' signal (e.g. the upper one in Fig.4, middle), then minimaly covering the remainder of the gadget can be done with only 9 balls. Similarly, if two or three wires are carrying a protrusion, a minimum covering of the remainder of the clause gadget has also cardinality 9. The case of three protrusions appears on the right in Fig.4, showing that even here 9 balls are still necessary to finish covering the gadget. Note that in general there may be several possible minimum coverings of the gadget, although only one is drawn here in each case.

According to these observations, it follows that the clause gadget can be minimaly covered by 10 balls if and only if no input protrusion is observed, in other words if and only if the corresponding clause is not satisfied. Otherwise, if at least one literal of the clause is set to 'True' (protrusion), meaning that the clause is satisfied, then only 9 balls are necessary to cover the remainder of the gadget.

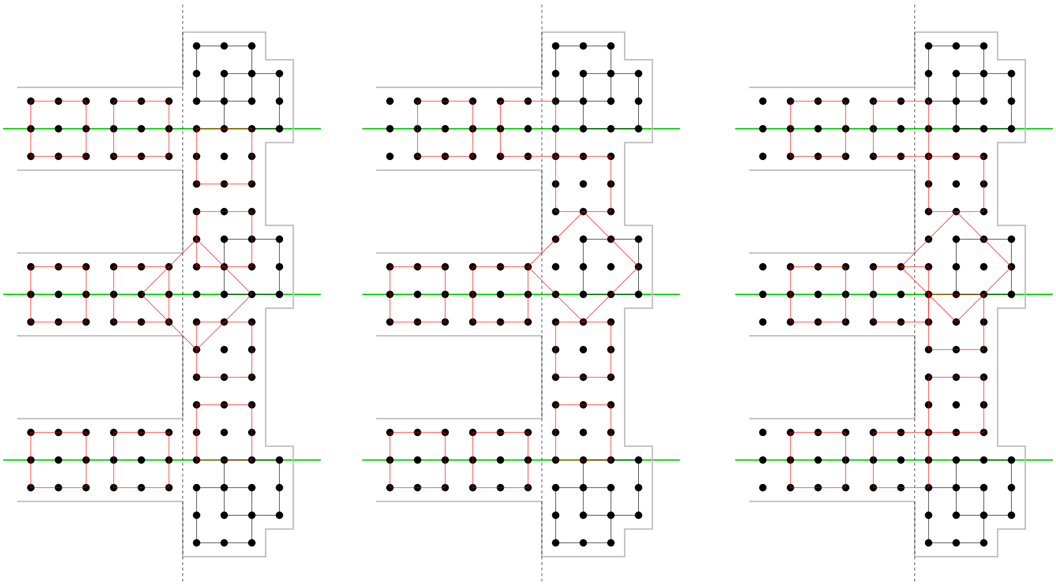

Fig. 4. Three minimum coverings of a clause gadget, depending on the following input signals (from left to right): False-False-False, True-False-False, True-True-True.

3.4 Overall Construction and Proof

Given a Planar-4 3-SAT formula $\phi(V, C)$, we are now ready to construct $\mathcal{S}(\phi)$ by drawing a variable gadget for each variable vertex in $G(\phi)$, a clause gadget for each clause vertex in $G(\phi)$, and drawing wires corresponding to the edges in $G(\phi)$, thus linking each literal (the extremity of a variable gadget) to every clause in which it occurs.

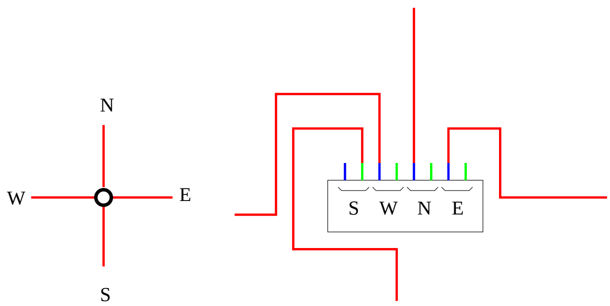

Fig. 5. Illustration of the transformation of a vertex of the planar orthogonal embedding into a variable gadget. In this case, the associated variable appears four times in ϕ, three times as a positive literal, and once as a negative literal.

Lemma 1 *The shape $\mathcal{S}(\phi)$ can be computed in polynomial time in the size of ϕ.*

PROOF. We know from [20] that every planar graph with n vertices (with degree ≤ 4) can be embedded in a rectilinear grid in polynomial time and space. This algorithm produces an orthogonal drawing such that edges are intersection free 4-connected discrete curves. Since our variable gadgets and clause gadgets have a constant size and our wires have constant width, and since ϕ is an instance of Planar-4 3-SAT , it is clear that the construction of $\mathcal{S}(\phi)$ can also be done in polynomial time and space. For example, Figure 5 illustrates how to bend the orthogonal drawing edges in order to connect them to our variable gadget extremities. □

In the following, let $w(\phi)$ denote the minimum number of balls necessary to cover the wires of $\mathcal{S}(\phi)$, and let $k(\phi(V,C)) = 72.|V| + w(\phi) + 9.|C|$.

Lemma 2 *If the formula ϕ is satisfiable, then there exists a covering of $\mathcal{S}(\phi)$ with $k(\phi)$ maximal balls.*

PROOF. Given a truth assignment T of the variables V of ϕ such that ϕ is satisfiable, the following algorithm builds a covering of $\mathcal{S}(\phi)$ with $k(\phi)$ maximal balls:

- cover the variable gadgets according to the truth assignment T ('True' or 'False' value for each variable): each one requires 72 balls allowing protrusions in each extremity carrying a 'True' assignement (Section 3.1);
- cover the wires: since the grid embedding of $G(\phi)$ is computed in polynomial time, so is $w(\phi)$; the protrusions from the extremities of the variables are transmitted to the clause gadgets;
- cover the clause gadgets: since ϕ is satisfiable, at least one input wire of each clause gadget carries a protrusion which implies that 9 maximal balls are enough to cover each clause gadgets (Section 3.3).

Altogether, $72.|V| + w(\phi) + 9.|C| = k(\phi)$ maximal balls are used in this covering. □

Lemma 3 *If there exist a covering of $\mathcal{S}(\phi)$ with $k(\phi)$ maximal balls, then the formula ϕ is satisfiable.*

PROOF. Suppose that there exists covering of $\mathcal{S}(\phi)$ with $k(\phi)$ maximal balls. By construction, $72.|V|$ plus $w(\phi)$ maximal balls are required to cover the $|V|$ variable gadgets and the wires of $\mathcal{S}(\phi)$. This leaves us with $k(\phi) - 72.|V| - w(\phi) = 9.|C|$ maximal balls to cover the clause gadgets. Since there are $|C|$ clause gadgets, each one is totally covered with 9 maximal balls in the covering, which is possible only if at least one input wire of each clause gadget carries a protrusion (Section 3.3). By construction, this means that the clauses are all satisfied, and in turn that ϕ is satisfiable. □

According to lemmas 2 and 3, there exists a truth assignement of the variables in V which satisfies all the clauses in ϕ if and only if there exists a covering of $\mathcal{S}(\phi)$ of cardinality $k(\phi) = 72.|V|+w(\phi)+9.|C|$. Thus, if any instance of the k-Medial Axis Problem could be solved in polynomial time, then we would have a polynomial time algorithm to solve the Planar-4 3-SAT Problem. Therefore, the k-MA Problem is NP-hard. It is also clear that k-MA problem is in NP, since we can easily verify in polynomial time that a set of k balls covers a discrete shape $\mathcal{S}$. Consequently, we have the following theorem:

Theorem 4 *k-MA is an NP-complete problem.*

As a consequence, finding a minimum subset of the medial axis of a shape $\mathcal{S}$, which still covers $\mathcal{S}$ is NP-hard.

4 Approximation Algorithms and Heuristics

Even if the theoretical problem is NP-hard, approximation and heuristics can be designed to reduce the cardinality of the discrete medial axis while keeping the reversibility. In the literature, several authors have discussed about simplification techniques to reduce the cardinality of the medial axis [5,15,6,13]. Dealing with NP-hard problems, we usually want to have bounded heuristics in the sense that the results given by the approximation algorithm will always be at a given distance from the optimal solution.

In the following, we first detail the simplification algorithm proposed by Ragnemalm and Borgefors [15]. Then, we compare their result with a simple but bounded heuristic derived from the MinSetCover problem.

4.1 Ragnemalm and Borgefors Simplification Algorithm

The algorithm is quite simple but provide interesting results: we first construct a covering map in which we count for each discrete point p, the number of discrete balls containing p. Basically, if a ball B contains a grid point p with a value of 1, then B is necessary to maintain the reconstruction and B belongs to any minimum medial axis. Based on this idea, the approximation algorithm can be sketched as follows: let $\mathcal{F} = \mathrm{MA}(\mathcal{S})$, for each ball B in $\mathcal{F}$ ordered by increasing radii, if all grid points of B have a value greater than or equal to 2, we remove B from $\mathcal{F}$ and decrease by one each grid point of B in the covering map.

The resulting set $\hat{\mathcal{F}}$ may be such that $|\hat{\mathcal{F}}| < |\mathcal{F}|$. In [15], the author illustrates the reduction rates with several shapes in dimension 2 but no formal simplification rate is given in the general case. In our experiments, instead of considering the medial axis of $\mathcal{S}$, we set $\mathcal{F} = \text{RMA}(\mathcal{S})$ [8].

If $\mathcal{F} = \{B_i, i = 1 \dots k\}$, the overall computational cost of this algorithm is $O(\sum_{i=1}^{k} |B_i| + k \log k)$.

4.2 Greedy Algorithm: a Bounded Heuristic

To have a bounded heuristic, let us consider another problem called the MinSetCover problem [9]: an instance $(\mathcal{S}, \mathcal{F})$ of the MinSetCover consists of a finite set $\mathcal{S}$ and a family $\mathcal{F}$ of subsets of $\mathcal{S}$, such that every element of $\mathcal{S}$ belongs to at least one subset of $\mathcal{F}$. The problem is to find the family of subsets $\mathcal{F}^*$ with minimum cardinality such that $\mathcal{F}^*$ still covers $\mathcal{S}$. From the optimization MinSetCover problem, we can define the following decision problem: can we cover $\mathcal{S}$ with a family $\mathcal{F}^*$ such that $|\mathcal{F}^*| \leq k$ for a given k ? This decision problem is known to be NP-complete [9]. Replacing $\mathcal{S}$ by a discrete object and $\mathcal{F}$ by the medial axis, we have a specific instance of the MinSetCover problem.

The greedy approximation algorithm is presented in 1. Even if this algorithm is simple, it provides a bounded approximation: if we denote $H(d) = \sum_{i=1}^{d} \frac{1}{i}$, $H_{\mathcal{F}} = H(\max |S|, S \in \mathcal{F})$ and $\mathcal{F}^*$ the minimum medial axis, the greedy algorithm produces a set $\hat{\mathcal{F}}$ such that:

$$|\hat{\mathcal{F}}| \leq H_{\mathcal{F}} \cdot |\mathcal{F}^*|$$

Algorithm 1: Greedy algorithm for MinSetCover.

Data: $\mathcal{S}$ and $\mathcal{F}$
Result: the approximated solution $\hat{\mathcal{F}}$
U $= \mathcal{S}$;
$\hat{\mathcal{F}} = \emptyset$;
while $U \neq \emptyset$ **do**
    Select $S \in \mathcal{F}$ that maximizes $|S \cap U|$;
    $U = U - S$;
    $\hat{\mathcal{F}} = \hat{\mathcal{F}} \cup \{S\}$;
return $\hat{\mathcal{F}}$

If we consider $\mathcal{S}$ as a discrete object and $\mathcal{F}$ given by the medial axis extraction, the medial axis simplification problem is a sub-problem of MinSetCover. Hence, Algorithm 1 provides a bounded heuristic for the medial axis reduction. Even if the bound is large according to experiments (see Section 4.3), this is at the time of writing the only known approximation algorithm for the minimum medial axis, for which we have an approximation factor. Despite the fact that Algorithm 1 has a computational cost in $O(|\mathcal{S}||\mathcal{F}| \min(|\mathcal{S}|, |\mathcal{F}|))$, a linear in time algorithm can be designed, *i.e.* in $O(\sum_{i=1}^{k} |B_i|)$. Yet, the implementation requires a bit more complicated data structure because instead of a covering map with numbers, we need to store a list of MA balls for each grid point.

4.3 Experiments

In Figure 6, we present some experiments of both approximation algorithms. Two observations can be addressed: first, the reduction rate is very interesting since almost half of the medial axis balls can be removed. Secondly, the computational time of both algorithms are almost similar.

Despite the fact that Ragnemalm and Borgefors's algorithm gives slightly better results, the theoretical bound provided by the greedy algorithm makes this approach a bit more satisfactory.

5 Discussion and Conclusion

In this paper, we prove that finding the medial axis of minimum cardinality of a discrete shape is a NP-hard problem. To do so, we provide a polynomial reduction from the Planar-4 3-SAT problem to the minimum medial axis problem. We also experimentally compare the output given by the greedy approximation algorithm with existing simplification algorithms.

In the proof, we have considered the Euclidean distance based medial axis. To derive a proof for the other metrics, new gadgets must be defined. Some cases are trivial, such as the d_8 case for which only the variable gadget must be redefined (see Figure 7), others may be difficult but in our point, the result may be the same.

Future works concern both the complexity of specific restrictions of the minimum medial axis problem, and the approximation algorithms. Concerning the theoretical part, the result we give induces the construction of very specific discrete shapes, whose genus depends on the number of cycles in the Planar-4 3-SAT instance. Thus, an important question is whether k-MA is still NP-complete in the case of connected discrete shapes without holes. As regards approximation algorithms, experiments show that the results of the greedy approximation algorithm are slightly worse than other existing algorithms. An important future work is to merge the two approaches to improve the results while keeping the bounded approximation.

Objet	$\mathcal{F} = \text{MA}(\mathcal{S})$	$\hat{\mathcal{F}}$ Ragnemalm et al.	$\hat{\mathcal{F}}$ Greedy
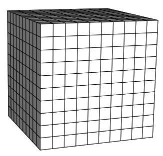	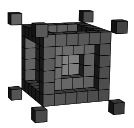 104	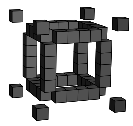 56 (-46%) [<0.01s]	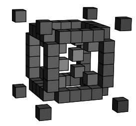 66 (-36%) [< 0.01s]
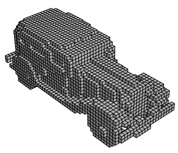	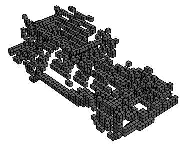 1292	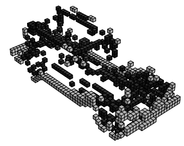 795 (-38%) [0.1s]	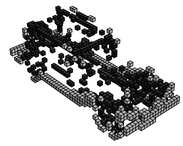 820 (-36%) [0.19s]
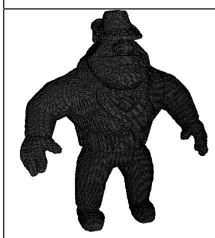	 17238	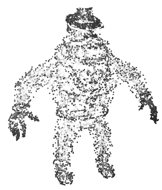 6177 (-64%) [48.53s]	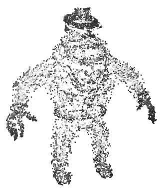 6553 (-62%) [57.79s]

Fig. 6. Experimental analysis of simplification algorithms: *(from left to right)* Discrete 3-D objects, the discrete medial axis, simplification obtained by [15], simplification obtained by the proposed greedy algorithm. The cardinality of the sets are given below the figure with the reduction ratio (in percent) and the computational time.

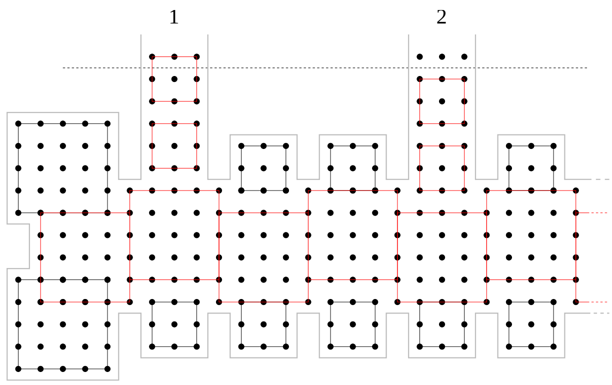

Fig. 7. Outline of a variable gadget for d_8

References

[1] L. J. Aupperle, H. E. Conn, J. M. Keil, and J. O'Rourke. Covering orthogonal polygons with squares. In *Proceedings of the 26th Annual Allerton Conf. Comm.*

Control Comput., pages 97–106, 1988.

[2] H. Blum. A transformation for extracting new descriptors of shape. In W. Wathen-Dunn, editor, *Models for the Perception of Speech and Visual Form*, pages 362–380, Cambridge, 1967. MIT Press.

[3] G. Borgefors. Distance transformations in arbitrary dimensions. *Computer Vision, Graphics and Image Processing*, 27:321–345, 1984.

[4] G. Borgefors. Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing*, 34(3):344–371, June 1986.

[5] G. Borgefors and I. Nyström. Quantitative shape analysis of volume images – reducing the set of centres of maximal spheres. In *Proc. SSAB Symposium on Image Analysis*, pages 5–8, Linköping, Sweden, March 1995.

[6] G. Borgefors and I. Nyström. Efficient shape representation by minimizing the set of centers of maximal discs/spheres. *Pattern Recognition Letters*, 18:465–472, 1997.

[7] H. Breu, J. Gil, D. Kirkpatrick, and M. Werman. Linear time euclidean distance transform algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(5):529–533, 1995.

[8] D. Coeurjolly and A. Montanvert. 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, mar 2007.

[9] T. Cormen, C. Leiserson, and R. Rivest. *Introduction à l'algorithmique*. Dunod, 1990.

[10] J. E. Goodman and J. O'Rourke, editors. *Handbook of Discrete and Computational Geometry*. CRC Press, 1997.

[11] T. Hirata. A unified linear-time algorithm for computing distance maps. *Information Processing Letters*, 58(3):129–133, May 1996.

[12] K. Jansen and H. Müller. The minimum broadcast time problem for several processor networks. *Theoretical Computer Science*, 147(1–2):69–85, 7 August 1995.

[13] F. Nilsson and P.-E. Danielsson. Finding the minimal set of maximum disks for binary objects. *Graphical models and image processing*, 59(1):55–60, January 1997.

[14] J.L. Pfaltz and A. Rosenfeld. Computer representation of planar regions by their skeletons. *Comm. of ACM*, 10:119–125, feb 1967.

[15] I. Ragnemalm and G. Borgefors. *The Euclidean Distance Transform,*, chapter Towards a minimal shape representation using maximal discs, pages 245–260. Linköping Studies in Science and Technology. Dissertations No. 304., Linköping University, apr 1993.

[16] E. Remy and E. Thiel. Medial axis for chamfer distances: computing look-up tables and neighbourhoods in 2D or 3D. *Pattern Recognition Letters*, 23(6):649–661, 2002.

[17] E. Remy and E. Thiel. Exact Medial Axis with Euclidean Distance. *Image and Vision Computing*, 23(2):167–175, 2005.

[18] A. Rosenfeld and J. L. Pfaltz. Sequential operations in digital picture processing. *Journal of the ACM*, 13(4):471–494, October 1966.

[19] T. Saito and J.-I. Toriwaki. Reverse distance transformation and skeletons based upon the euclidean metric for n-dimensionnal digital pictures. *IECE Trans. Inf. & Syst.*, E77-D(9):1005–1016, September 1994.

[20] R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Comput.*, pages 421–444, 1987.