

Surface thinning in 3D cubical complexes

John Chaussard and Michel Couprie

Université Paris-Est, Laboratoire d'Informatique Gaspard-Monge, Équipe A3SI,
ESIEE Paris, France

Abstract. We introduce a parallel thinning algorithm with directional substeps based on the collapse operation, which is guaranteed to preserve topology and to provide a thin result. Then, we propose two variants of a surface-preserving thinning scheme, based on this parallel directional thinning algorithm. Finally, we propose a methodology to produce filtered surface skeletons, based on the above thinning methods and the recently introduced discrete λ -medial axis.

1 Introduction

The notion of skeleton plays a major role in shape analysis and recognition. Informally, the skeleton of an object X is a subset of X which is 1) thin, 2) centered in X and 3) topologically equivalent¹ to X . To extract skeletons from geometric objects, different methods have been proposed, relying on different frameworks: discrete geometry [8, 13, 19, 22, 14], digital topology [12, 28, 27, 21], mathematical morphology [24, 26], computational geometry [2, 3, 20], and partial differential equations [25].

In this paper, we focus on skeletons in the discrete 3-dimensional cubical space. In such spaces, topology preservation is usually guaranteed by the use of topology-preserving thinning methods, based *e.g.* on simple point deletion [15, 11]. The centering of the skeleton may be achieved to some extent by the use of parallel thinning methods [6] or by imposing geometrical constraints such as the preservation of the centers of maximal included balls [12]. However it is generally difficult, in discrete spaces, to satisfy conditions 1), 2) and 3) together. In particular, the methods evoked above fail to guarantee a thin result.

In order to overcome these limitations, we adopt in this article the framework of cubical complexes. Abstract (cubical) complexes have been promoted in particular by V. Kovalevsky [17] in order to provide a sound topological basis for image analysis. Intuitively, a cubical complex may be thought of as a set of elements having various dimensions (*e.g.* cubes, squares, edges, vertices) glued together according to certain rules (see Fig. 3). In this framework, we can say that a complex is “thin” if it does not contain any 3-dimensional element (cube). The thinning methods that will be proposed in the sequel of this paper produce thin results in this sense.

¹ To be more precise, we say that a transformation Ψ “preserves topology” if X is homotopy-equivalent to $\Psi(X)$ for any X .

In the framework of cubical complexes, the notion of critical kernel introduced by G. Bertrand constitutes a powerful tool to study parallel homotopic thinning in any dimension, which unifies and encompasses previous works on parallel thinning [4, 7]. Indeed, the very notion of critical kernel may be seen as thinning scheme, which consists of iteratively computing the critical kernel of the result of the previous step. Critical kernels may also be used to design new algorithms, as well as to check the topological validity of existing ones [5, 6]. However, thinning algorithms based on critical kernels do not guarantee that the result has a dimension strictly lower than the input. Early work on thinning in cubical complexes with the aim of reducing dimension can be found in [18].

In our work, topology preservation will be ensured by the use of the *collapse operation*, which is an elementary topology-preserving transformation that has been introduced by J.H.C. Whitehead [29] and plays an important role in combinatorial topology. It can be seen as a discrete analogue of a retraction, that is, a continuous deformation of an object onto itself.

The contributions of this article are the following. We introduce a parallel thinning algorithm with directional substeps based on collapse, which is guaranteed to preserve topology and to provide a thin result. Then, we propose two variants of a surface-preserving thinning scheme, based on this parallel directional thinning algorithm. Finally, we propose a methodology to produce filtered surface skeletons, based on the above thinning methods and the recently introduced discrete λ -medial axis [9].

2 Basic notions

2.1 Cubical complexes

Let \mathbb{Z} be the set of integers, we consider the family of sets \mathbb{F}_0^1 and \mathbb{F}_1^1 , such that $\mathbb{F}_0^1 = \{\{a\} \mid a \in \mathbb{Z}\}$ and $\mathbb{F}_1^1 = \{\{a, a + 1\} \mid a \in \mathbb{Z}\}$. Any subset f of \mathbb{Z}^n such that f is the cartesian product of m elements of \mathbb{F}_1^1 and $(n - m)$ elements of \mathbb{F}_0^1 is called a face or an m -face of \mathbb{Z}^n , m is the dimension of f , we write $\dim(f) = m$. A 0-face is called a *vertex*, a 1-face is an *edge*, a 2-face is a *square*, and a 3-face is a *cube* (see Fig. 1).




Fig. 1. Graphical representations of: (a) a 0-face, (b) a 1-face, (c) a 2-face, (d) a 3-face.

We denote by \mathbb{F}^n the set composed of all faces in \mathbb{Z}^n .

Let $f \in \mathbb{F}^n$. We set $\hat{f} = \{g \in \mathbb{F}^n \mid g \subseteq f\}$, and $\hat{f}^* = \hat{f} \setminus \{f\}$. Any element of \hat{f} is a *face* of f , and any element of \hat{f}^* is a *proper face* of f . We call *star* of f the

set $\check{f} = \{g \in \mathbb{F}^n \mid f \subseteq g\}$, and we write $\check{f}^* = \check{f} \setminus \{f\}$: any element of \check{f} is a *coface* of f . It is plain that $g \in \check{f}$ iff $f \in \check{g}$.

A set X of faces in \mathbb{F}^n is a *cell*, or *m-cell*, if there exists an m -face $f \in X$ such that $X = \hat{f}$. The *closure* of a set of faces X is the set $X^- = \cup\{\hat{f} \mid f \in X\}$. The set \overline{X} is $\mathbb{F}^n \setminus X$.

A finite set X of faces in \mathbb{F}^n is a cubical *complex* if $X = X^-$, and we write $X \preceq \mathbb{F}^n$. Any subset Y of X which is also a complex is a *subcomplex* of X , and we write $Y \preceq X$.

A face $f \in X$ is a *facet* of X if f is not a proper face of any face of X . We denote by X^+ the set composed of all facets of X . A complex X is *pure* if all its facets have the same dimension. The *dimension* of X is $\dim(X) = \max\{\dim(f) \mid f \in X\}$. If $\dim(X) = d$, then we say that X is a d -complex.

2.2 From binary images to complexes

Traditionally, a binary image (see Fig. 2a) is represented as a subset of \mathbb{Z}^2 (or \mathbb{Z}^3 for 3D images). Let $S \subseteq \mathbb{Z}^n$, the elements of S represent the pixels (or voxels) of the image, often called object pixels or black pixels. The set $\overline{S} = \mathbb{Z}^n \setminus S$ represents the background (set of white pixels or voxels).




Fig. 2. (a): A binary image, represented by a set $S \subseteq \mathbb{Z}^2$ (black pixels). (b): The complex $X = \Phi(S)^-$.

To obtain a complex from a subset of \mathbb{Z}^n , we consider the following bijection which associates, to each point of \mathbb{Z}^n , a facet of \mathbb{F}^n . Let $x = (x_1, \dots, x_n) \in \mathbb{Z}^n$, define $\Phi(x) = \{x_1, x_1 + 1\} \times \dots \times \{x_n, x_n + 1\}$. See Fig. 2 an illustration in 2D, where the image of each pixel by Φ is a 2-face (a square). For example, the black pixel x in the first row has coordinates $(3, 0)$; the corresponding 2-face is $\Phi(x) = \{3, 4\} \times \{0, 1\} = \{(3, 0), (3, 1), (4, 0), (4, 1)\}$. The map Φ is straightforwardly extended to sets: $\Phi(S) = \{\Phi(x) \mid x \in S\}$. Then, the complex that we associate to a set S is $\Phi(S)^-$ (See Fig. 2b). Notice that the datum of an n D binary image is equivalent to the datum of a pure n -complex in \mathbb{F}^n .

2.3 Collapse

The collapse operation consists of removing two distinct elements (f, g) from a complex X under the condition that g is contained in f and is not contained

in any other element of X . This operation may be repeated several times (see Fig. 3). A more precise definition follows.




Fig. 3. (a-e): A series of complexes that illustrate a sequence of collapse operations.

Let $X \preceq \mathbb{F}^n$, and let f, g be two faces of X . The face g is *free for X* , and the pair (f, g) is a *free pair for X* if f is the only face of X such that g is a proper face of f . In other terms, (f, g) is a free pair for X whenever $\check{g}^* \cap X = \{f\}$. In this case, we say that the complex $X \setminus \{f, g\}$ is an *elementary collapse of X* . It can be easily seen that if (f, g) is a free pair for X and $\dim(f) = m$, then f is a facet and $\dim(g) = (m - 1)$.

Let $X \preceq \mathbb{F}^n$ and $Y \preceq \mathbb{F}^n$ be two complexes. We say that X *collapses onto Y* if there exists a sequence of complexes (X_0, \dots, X_ℓ) of \mathbb{F}^n such that $X = X_0$, $Y = X_\ell$ and for all $i \in \{1, \dots, \ell\}$, X_i is an elementary collapse of X_{i-1} (see Fig. 3).

Let f_0, f_ℓ be two n -faces of \mathbb{F}^n . An $(n - 1)$ -*path from f_0 to f_ℓ* is a sequence $\pi = (f_0, \dots, f_\ell)$ of faces of \mathbb{F}^n such that for all $i \in \{0, \dots, \ell\}$, either i is even and f_i is an n -face, or i is odd and f_i is an $(n - 1)$ -face with $\check{f}_i^* = \{f_{i-1}, f_{i+1}\}$.

The following proposition will serve us to prove the thinness of our skeletons.

Proposition 1. *Let $X \preceq \mathbb{F}^n$ be an n -complex, with $n > 0$. Then X has at least one free $(n - 1)$ -face.*

Proof. Since X is an n -complex (hence X is finite) there exists an n -face a in X and an n -face b in \overline{X} . Obviously, there exists an $(n - 1)$ -path from a to b . Let h be the first n -face of π that is not in X , let k be the last n -face of π before h (thus k is in X), and let $e = k \cap h$ be the $(n - 1)$ -face of π between k and h . Since k and h are the only two n -faces of \mathbb{F}^n that contain e , we see that the pair (k, e) is free for X . \square

3 Parallel directional thinning based on collapse

The most “natural” way to thin an object consists of removing some of its border elements in parallel, in a symmetrical manner. However, parallel deletion of free pairs does not, in general, guarantee topology preservation: see for example

Fig. 4, where simultaneously removing all free pairs would split the object into separate components.




Fig. 4. (a): A 2-complex X , each free pair (f_i, g_i) is symbolized by an arrow from g_i to f_i . (b): The complex obtained by removing simultaneously from X all the pairs that are free for X .

In the framework of 2D digital topology, a popular method due to A. Rosenfeld [23] consists of dividing each thinning step into directional substeps. In each substep, only simple points that have no neighbor belonging to the object in one of the four main directions (north, south, east, west) are candidates for deletion.

However, this method cannot be straightforwardly extended to 3D. In fact, the question of knowing whether this strategy has a “natural extension” to 3D was cited among three open questions relative to digital topology by Kong, Litherland and Rosenfeld in [16] (question 547). An answer to this question was recently given in [7], based on the critical kernels framework. Here, we show that the directional strategy can be directly adapted in the framework of cubical complexes, and provides interesting results.

First, we need to define the *direction* and the *orientation* of a free face.

Let $f \in \mathbb{F}^n$, the *center of f* is the center of mass of the points in f , that is, $c_f = \frac{1}{|f|} \sum_{a \in f} a$. The center of f is an element of $[\frac{\mathbb{Z}}{2}]^n$, where $\frac{\mathbb{Z}}{2}$ denotes the set of half integers. Let $X \preceq \mathbb{F}^n$, let (f, g) be a free pair for X , and let c_f and c_g be the respective centers of the faces f and g . We denote by $V(f, g)$ the vector $(c_f - c_g)$ of $[\frac{\mathbb{Z}}{2}]^n$.

Let (f, g) be a free pair, the vector $V(f, g)$ has only one non-null coordinate. We define $Dir(f, g)$ as the index of the non-null coordinate of $V(f, g)$. Thus, $Dir(\cdot)$ is a surjective function from $\mathbb{F}^n \times \mathbb{F}^n$ to $\{1, \dots, n\}$ such that, for all free pairs (f, g) and (i, j) for X , $Dir(f, g) = Dir(i, j)$ if and only if $V(f, g)$ and $V(i, j)$ are collinear. The number $Dir(f, g)$ is called the *direction* of the free pair (f, g) . The free pair (f, g) has a *positive orientation*, and we write $Orient(f, g) = 1$, if the non-null coordinate of $V(f, g)$ is positive; otherwise (f, g) has a *negative orientation*, and we write $Orient(f, g) = 0$.

Let us now state an elementary property of collapse, which gives a necessary and sufficient condition under which two collapse operations may be performed in parallel while preserving topology.

Proposition 2. *Let $X \preceq \mathbb{F}^n$, and let (f, g) and (k, ℓ) be two distinct free pairs for X . The complex X collapses onto $X \setminus \{f, g, k, \ell\}$ if and only if $f \neq k$.*

Proof. If $f = k$, then it is plain that (k, ℓ) is not a free pair for $Y = X \setminus \{f, g\}$ as $k = f \notin Y$. Also, (f, g) is not free for $X \setminus \{k, \ell\}$. If $f \neq k$, then we have $g \neq \ell$,

$\check{g}^* \cap X = \{f\}$ (g is free for X) and $\check{\ell}^* \cap X = \{k\}$ (ℓ is free for X). Thus, we have $\check{\ell}^* \cap Y = \{k\}$ as $\ell \neq g$ and $k \neq f$. Therefore, (k, ℓ) is a free pair for Y . \square

From Prop. 2, the following corollary is immediate.

Corollary 3. *Let $X \preceq \mathbb{F}^n$, and let $(f_1, g_1), \dots, (f_m, g_m)$ be m distinct free pairs for X such that, for all $a, b \in \{1, \dots, m\}$ (with $a \neq b$), $f_a \neq f_b$. The complex X collapses onto $X \setminus \{f_1, g_1, \dots, f_m, g_m\}$.*

Considering two distinct free pairs (f, g) and (i, j) for $X \preceq \mathbb{F}^n$ such that $Dir(f, g) = Dir(i, j)$ and $Orient(f, g) = Orient(i, j)$, we have $f \neq i$. From this observation and Cor. 3, we deduce the following property.

Corollary 4. *Let $X \preceq \mathbb{F}^n$, and let $(f_1, g_1), \dots, (f_m, g_m)$ be m distinct free pairs for X having all the same direction and the same orientation. The complex X collapses onto $X \setminus \{f_1, g_1, \dots, f_m, g_m\}$.*

We say that a d -face of X is a *border face* if it contains a free $(d-1)$ -face. Define $Border(X)$ as the set of all border faces of X . We are now ready to introduce our directional thinning algorithm (Alg. 1).

Algorithm 1: *ParDirCollapse*(X, W, ℓ)

Data: A cubical complex $X \preceq \mathbb{F}^n$, a subset $W \subseteq X$ that contains faces of X which should not be removed, and $\ell \in \mathbb{N}$, the number of layers of free faces which should be removed from X

```

1 while there exists free pairs for  $X$  in  $X \setminus W$  and  $\ell > 0$  do
2    $L = Border(X)^-$ ;
3   for  $t = 1 \rightarrow n$  do
4     for  $s = 0 \rightarrow 1$  do
5       for  $d = n \rightarrow 1$  do
6          $E = \{(f, g) \text{ free for } X \mid g \notin W, f \notin W,$ 
7            $Dir(f, g) = t, Orient(f, g) = s, \dim(f) = d\};$ 
8          $G = \{(f, g) \in E \mid f \in L \text{ and } g \in L\};$ 
9          $X = X \setminus G;$ 
9    $\ell = \ell - 1;$ 
10 return  $X;$ 

```

Let us first comment on the notion of “layer” and the role of the sets L and G (lines 2 and 7). Intuitively, we want a single execution of the loop line 1 to only affect facets that were on the border of X at the beginning of the loop, in order to ensure that the thinning will not reduce certain parts more quickly than others due to the scanning of directions. Thus, by line 7 we ensure that the layer which is removed in one iteration of the loop line 1 is included in L , which contains all the border facets of the complex X and all the faces included in those.

Remark that different definitions of the functions $Dir()$ and $Orient()$ could be given, corresponding to different orders in which directions and orientations are scanned in Alg. 1, and yielding different results. However in general some arbitrary choices must be made in order to thin a 3-complex into a 2-complex. The choice of these two functions is the only arbitrary choice to be made, after which the results are uniquely defined.

For any complex $X \preceq \mathbb{F}^n$, any subset $W \subseteq X$ and for any $\ell \in \mathbb{N}$, it follows from Cor. 4 that X collapses onto $ParDirCollapse(X, W, \ell)$. Furthermore, if W does not contain any n -face nor any $(n - 1)$ -face that is not a facet of X , and if $\ell = +\infty$, then by Prop. 1 it can easily be deduced that $ParDirCollapse(X, W, \ell)$ contains no n -face.

Notice that checking whether a face is free or not, is quite easy to implement and can be done in constant time, whatever the dimension of the face, if n is considered as fixed. Indeed, Alg. 1 can be used to thin a cubical complex of any dimension. When a free pair (f, g) is removed from X , it is sufficient to scan the faces of f in order to check the appearance of new free faces, avoiding to search the whole complex for free faces. Thus, Alg. 1 may be easily implemented to run in linear time complexity (proportionally to the number of faces of the complex).

As shown in Fig.5, when the input complex has a nearly constant thickness, it is possible to obtain a surface skeleton of X by choosing a convenient value of ℓ as the last parameter of Alg. 1. However, in most cases, it is not possible to find a value of ℓ which is satisfying for the whole complex. In the next section, we explain how to use the directional thinning strategy to obtain a surface skeleton from a complex without having to tune any “thickness” parameter.




Fig. 5. (a): A 3-complex X having the shape of a thick tube with a skewed axis. (b): The result of $ParDirCollapse(X, \emptyset, 10)$ is not thin, some 3-faces remain in the object. (c): The result of $ParDirCollapse(X, \emptyset, 15)$ is a 2-complex. (d): The result of $ParDirCollapse(X, \emptyset, \infty)$ does not contain any cube nor any square, it is composed of edges and points.

4 Detection of surfaces

In this section, two algorithms are proposed in order to perform a surface-preserving thinning of a cubical complex. Both algorithms are based on the same principle: a single layer of faces is removed using Alg. 1, then some 2-faces are detected and kept safe from any future removal. This process is repeated until stability.

For Alg. 2 (*CollapseSurface*), the detected 2-faces are simply those that are not contained in any 3-face of the current complex, *i.e.*, the 2-facets of X .

Algorithm 2: *CollapseSurface*(X, W)

Data: A three-dimensional cubical complex $X \preceq \mathbb{F}^3$, a subset $W \subseteq X$ that contains faces of X which should not be removed

- 1 **while** *there exists free pairs for X in $X \setminus W$* **do**
- 2 $\left[\begin{array}{l} \text{ParDirCollapse}(X, W, 1); \\ W = W \cup \{f \in X^+ \mid \dim(f) = 2\}; \end{array} \right.$
- 3
- 4 **return** X ;

Fig. 6a shows the result of *CollapseSurface*(X, \emptyset), where X is the tube shown in Fig. 5a. The resulting complex is a 2-complex containing some “branches” and “surface patches” (as shown on the detailed view) which do not represent significant “surfacic features” of the original object. This motivates the introduction of a variant of this method, which imposes a more restrictive condition to preserve 2-facets from removal.

We say that a 2-facet $f \in X^+$ is a *2d-isthmus* of X if it contains no free edge, in other words, if each edge in f is included in a facet of X distinct from f . In Alg. 3 (*CollapseIsthmus*), the 2-facets that are detected and kept safe from further removal are the 2d-isthmuses of X .

Algorithm 3: *CollapseIsthmus*(X, W)

Data: A three-dimensional cubical complex $X \preceq \mathbb{F}^3$, a subset $W \subseteq X$ that contains faces of X which should not be removed

- 1 **while** *there exists free pairs for X in $X \setminus W$* **do**
- 2 $\left[\begin{array}{l} \text{ParDirCollapse}(X, W, 1); \\ W = W \cup \{f \in X \mid f \text{ is a 2d-isthmus of } X\}; \end{array} \right.$
- 3
- 4 **return** X ;

Notice that, following Prop. 1, the results of Alg. 2 and Alg. 3 are thin (they do not contain any 3-face), provided that the input parameter W is set to \emptyset . However, thanks to parameter W , it is possible to constrain these algorithms to preserve selected parts of the original object, in addition to the surface parts that are automatically detected.

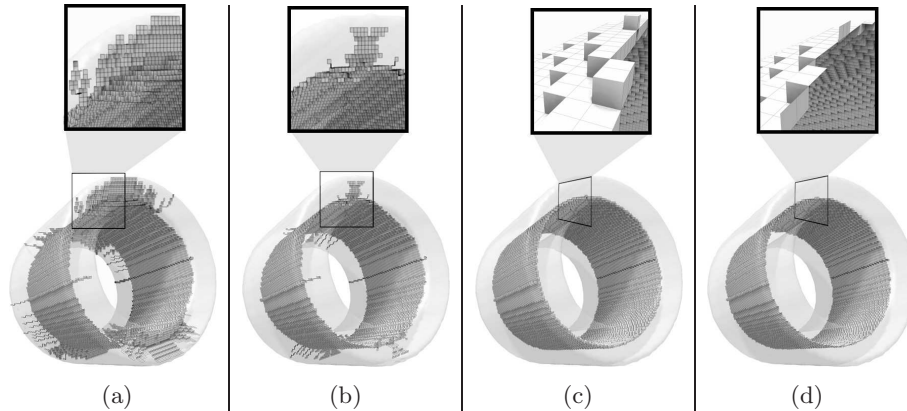


Fig. 6. Results of different algorithms performed on the complex X presented in Fig 5a. (a) : $CollapseSurface(X, \emptyset)$. (b) : $CollapseIsthmus(X, \emptyset)$. (c) : The discrete λ -medial axis (Sec. 5) of X , with $\lambda = 18$. (d) : Result of the method described in Sec. 5.3 performed on X , with $\lambda = 18$.

Fig. 6b shows the result of $CollapseIsthmus(X, \emptyset)$, where X is the skewed tube shown in Fig. 5a. The resulting complex contains less branches and spurious surfaces (as shown on the detailed view) than the complex obtained with $CollapseSurface(X, \emptyset)$.

The results obtained show that both algorithms presented above allows one to obtain a 2-dimensional skeleton from a 3-complex, containing important “shape information” from the initial object. However, even if Alg. 3 produces better results than Alg. 2, it fails in obtaining a completely satisfactory skeleton. In the next section, we introduce a methodology which includes a filtering step based on the Euclidean distance, providing a better robustness to contour irregularities and a lower sensitivity to the orientation of shape features in the space.

5 Skeleton filtering

A major difficulty when using the skeleton in applications (*e.g.*, shape recognition, shape analysis), is its sensitivity to small contour perturbations, in other words, its lack of stability. A recent survey [1] summarizes selected relevant studies dealing with this topic. This difficulty can be expressed mathematically: the transformation which associates a shape to its skeleton (or medial axis) is only semi-continuous. This fact, among others, explains why it is usually necessary to add a filtering step (or pruning step) to any method that aims at computing a skeleton. Hence, there is a rich literature devoted to skeleton pruning, in which different criteria were proposed in order to discard “spurious” skeleton points or branches.

In 2005, Chazal and Lieutier [10] introduced the λ -medial axis and studied its properties, in particular those related to stability. A major outcome of [10]

is the following property: informally, for “regular” values of λ , the λ -medial axis of X remains stable under perturbations of \overline{X} that are small with regard to the Hausdorff distance. Typical non-regular values are radii of locally largest maximal balls.

5.1 Intuitive presentation of the λ -medial axis

Consider a bounded subset X of \mathbb{R}^n , as for example, for $n = 2$, the region enclosed by the solid curve depicted in Fig. 7 (left). The *medial axis of X* consists of the points of X that have several closest points on the boundary of X . For example in Fig. 7, the boundary points closest to x are a and b , the boundary points closest to x' are a' and b' , and the only boundary point closest to x'' is a'' .

Let λ be a non-negative real number, the λ -*medial axis of X* is the set of points x of X such that the smallest ball including all boundary points that are closest to x , has a radius greater than or equal to λ . Notice that the 0-medial axis of X is equal to X , and that any λ -medial axis with $\lambda > 0$ is included in the medial axis. We show in Fig. 7 (right) two λ -medial axes with different values of λ .




Fig. 7. Illustration of the λ -medial axis. Left: Points x, x' and x'' and their respective closest boundary points. Top right: λ -medial axis with $\lambda = \epsilon$, a very small positive real number. Bottom right: λ -medial axis with $\lambda = d(a', b') + \epsilon$, where $d(a', b')$ is the distance between a' and b' .

Intuitively, in 2D, every “bump” on the boundary of the object generates a medial axis branch. The goal of the filtering is to eliminate those branches which are due to insignificant contour irregularities (noise). Roughly speaking, in the λ -medial axis, the branches that remain correspond to contour features having a “width” greater than the parameter value λ .

5.2 Definition of the discrete λ -medial axis

The original definition of the λ -medial axis holds and make sense in the continuous Euclidean n -dimensional space. In [9], we introduced the definition of a discrete λ -medial axis (DLMA) in \mathbb{Z}^n . We evaluated experimentally its stability and rotation invariance. Furthermore, we introduced a variant of the DLMA

which may be computed in linear time, for which the results are very close to those of the DLMA.

Notice that the DLMA applies on a binary image (*i.e.*, a set of voxels or a subset of \mathbb{Z}^3), not on a complex.

Let $x, y \in \mathbb{R}^n$, we denote by $d(x, y)$ the Euclidean distance between x and y , in other words, $d(x, y) = (\sum_{k=1}^n (y_k - x_k)^2)^{\frac{1}{2}}$. Let $S \subseteq \mathbb{R}^n$, we set $d(y, S) = \min_{x \in S} \{d(y, x)\}$.

Let $x \in \mathbb{R}^n, r \in \mathbb{R}^+$, we denote by $B_r(x)$ the *ball of radius r centered on x* , defined by $B_r(x) = \{y \in \mathbb{R}^n \mid d(x, y) \leq r\}$.

For each point $x \in \mathbb{Z}^n$, we define the *direct neighborhood of x* as $N(x) = \{y \in \mathbb{Z}^n \mid d(x, y) \leq 1\}$. The direct neighborhood comprises $2n + 1$ points.

Let S be a nonempty subset of \mathbb{Z}^n , and let $x \in \mathbb{Z}^n$. The *projection of x on S* , denoted by $\Pi_S(x)$, is the set of points y of S which are at minimal distance from x ; more precisely,

$$\Pi_S(x) = \{y \in S \mid \forall z \in S, d(y, x) \leq d(z, x)\}.$$

Let S be a finite subset of \mathbb{Z}^n , we denote by $R(S)$ the radius of the smallest ball enclosing S , that is, $R(S) = \min\{r \in \mathbb{R} \mid \exists y \in \mathbb{R}^n, B_r(y) \supseteq S\}$.

Transposing directly the definition of the λ -medial axis to the discrete grid \mathbb{Z}^n would yield unsatisfactory results (see [9]), this is why we need the following notion. Let $S \subseteq \mathbb{Z}^n$, and let $x \in S$. The *extended projection of x on \overline{S}* , denoted by $\Pi_{\overline{S}}^e(x)$, is the union of the sets $\Pi_{\overline{S}}(y)$, for all y in $N(x)$ such that $d(y, \overline{S}) \leq d(x, \overline{S})$.

Let S be a finite subset of \mathbb{Z}^n , and let $\lambda \in \mathbb{R}^+$. We define the function \mathcal{F}_S which associates, to each point x of S , the value $\mathcal{F}_S(x) = R(\Pi_{\overline{S}}^e(x))$. The *discrete λ -medial axis of S* , denoted by $DLMA(S, \lambda)$, is the set of points x in S such that $\mathcal{F}_S(x) \geq \lambda$.

Fig. 6c shows the DLMA of the object presented in Fig. 5a, with $\lambda = 18$. For more details, illustrations and performance analysis, see [9].

5.3 Surface skeleton based on the DLMA

Let us now describe our methodology to obtain a thin filtered surface skeleton from a binary image (a set of voxels). Let S denote our original set, a finite subset of \mathbb{Z}^3 .

First, we compute the discrete λ -medial axis of S for a chosen value of λ . From here, we will consider the complex $X = \Phi(S)^-$ and the set $W = \Phi(DLMA(S, \lambda))$.

Notice that we have no guarantee that the complex W^- be topologically equivalent to the complex X (see counter-examples in [9]). Thus our next step consists of computing, using Alg. 1, the complex $Y = ParDirCollapse(X, W, +\infty)$ that contains W and that is topologically equivalent to X .

Now, as W is made of 3-faces (see the close-up of Fig. 6c), Y is a 3-complex and cannot be seen as thin in a strict sense. Then to achieve our initial goal we compute $Z = CollapseIsthmus(Y, \emptyset)$ which is both thin (following Prop. 1 there is no 3-face in this complex, see the close-up of Fig. 6d for an illustration) and topologically equivalent to X (by Cor. 4). The centering of Z in X is achieved

thanks to the use of the DLMA, based on the Euclidean distance. The parameter λ can be tuned in order to adjust the filtering to the characteristics (size, smoothness of contours . . .) of the input shape, and to the requirements of the user.

Figures 6d and 8 show various results using this methodology. It can be seen that the resulting 2-complexes indeed capture the main surfacic features of the original objects, without spurious branches or surface patches. A significant advantage of the 2D nature of the obtained skeletons, is to enable an easy analysis of important shape features such as intersections of surface parts.

6 Conclusion

We introduced in this work a methodology for obtaining filtered, thin surface skeletons of 3D objects. This methodology is based on a robust medial axis extraction method (DLMA), a parallel directional thinning algorithm based on collapse, and a surface detection strategy. The algorithms involved in this methodology can be implemented to run in linear time complexity.

Aknowledgements

This work has been partially supported by the “ANR BLAN07-2_184378 MicroFiss” project.

References

1. D. Attali, J.-D. Boissonnat, and H. Edelsbrunner. Stability and computation of the medial axis — a state-of-the-art report. In T. Möller, B. Hamann, and B. Russell, editors, *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, pages 1–19. Springer-Verlag, 2009. to appear.
2. D. Attali and J.O. Lachaud. Delaunay conforming iso-surface, skeleton extraction and noise removal. *Computational Geometry: Theory and Applications*, 19:175–189, 2001.
3. D. Attali and A. Montanvert. Modelling noise for a better simplification of skeletons. In *Procs. International Conference on Image Processing (ICIP)*, volume 3, pages 13–16, 1996.
4. G. Bertrand. On critical kernels. *Comptes Rendus de l’Académie des Sciences, Série Math.*, I(345):363–367, 2007.
5. G. Bertrand and M. Couprie. Two-dimensional parallel thinning algorithms based on critical kernels. *Journal of Mathematical Imaging and Vision*, 31(1):35–56, 2008.
6. Gilles Bertrand and Michel Couprie. A new 3D parallel thinning scheme based on critical kernels. In *Discrete Geometry for Computer Imagery*, volume 4245 of *Lecture Notes in Computer Science*, pages 580–591. Springer, 2006.
7. Gilles Bertrand and Michel Couprie. On parallel thinning algorithms: minimal non-simple sets, P-simple points and critical kernels. *Journal of Mathematical Imaging and Vision*, 35(1):23–35, 2009.

8. G. Borgefors, I. Ragnemalm, and G. Sanniti di Baja. The Euclidean distance transform: finding the local maxima and reconstructing the shape. In *Procs. of the 7th Scandinavian Conference on Image Analysis*, volume 2, pages 974–981, 1991.
9. John Chaussard, Michel Couprie, and Hugues Talbot. A discrete lambda-medial axis. In *15th Discrete Geometry for Computer Imagery (DGCI'09)*, Lecture Notes in Computer Science, pages 1–12, 2009. To appear.
10. F. Chazal and A. Lieutier. The lambda medial axis. *Graphical Models*, 67(4):304–331, 2005.
11. Michel Couprie and Gilles Bertrand. New characterizations of simple points in 2D, 3D and 4D discrete spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(4):637–648, April 2009.
12. E.R. Davies and A.P.N. Plummer. Thinning algorithms: a critique and a new methodology. *Pattern Recognition*, 14:53–63, 1981.
13. Y Ge and J.M. Fitzpatrick. On the generation of skeletons from discrete Euclidean distance maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(11):1055–1066, 1996.
14. Wim H. Hesselink and Jos B.T.M. Roerdink. Euclidean skeletons of digital image and volume data in linear time by the integer medial axis transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(12):2204–2217, 2008.
15. T. Yung Kong and A. Rosenfeld. Digital topology: introduction and survey. *Comp. Vision, Graphics and Image Proc.*, 48:357–393, 1989.
16. T.Y. Kong, R. Litherland, and A. Rosenfeld. Problems in the topology of binary digital images. In *Open problems in topology*, pages 376–385. Elsevier, 1990.
17. V.A. Kovalevsky. Finite topology as applied to image analysis. *Computer Vision, Graphics and Image Processing*, 46:141–161, 1989.
18. Lu Liu. 3d thinning on cell complexes for computing curve and surface skeletons. Master's thesis, Washington Mاستر in Saint Louis, may 2009.
19. G. Malandain and S. Fernández-Vidal. Euclidean skeletons. *Image and Vision Computing*, 16:317–327, 1998.
20. R.L. Ogniewicz and O. Kübler. Hierarchic Voronoi skeletons. *Pattern Recognition*, 28(33):343–359, 1995.
21. C. Pudney. Distance-ordered homotopic thinning: a skeletonization algorithm for 3D digital images. *Computer Vision and Image Understanding*, 72(3):404–413, 1998.
22. E. Rémy and E. Thiel. Exact medial axis with Euclidean distance. *Image and Vision Computing*, 23(2):167–175, 2005.
23. A. Rosenfeld. A characterization of parallel thinning algorithms. *Information and Control*, 29:286–291, 1975.
24. J. Serra. *Image analysis and mathematical morphology*. Academic Press, 1982.
25. K. Siddiqi, S. Bouix, A. Tannenbaum, and S. Zucker. The Hamilton-Jacobi skeleton. In *International Conference on Computer Vision (ICCV)*, pages 828–834, 1999.
26. P. Soille. *Morphological image analysis*. Springer-Verlag, 1999.
27. H. Talbot and L. Vincent. Euclidean skeletons and conditional bisectors. In *Procs. VCIP'92, SPIE*, volume 1818, pages 862–876, 1992.
28. L. Vincent. Efficient computation of various types of skeletons. In *Procs. Medical Imaging V, SPIE*, volume 1445, pages 297–311, 1991.
29. J.H.C. Whitehead. Simplicial spaces, nuclei and m -groups. *Proceedings of the London Mathematical Society*, 45(2):243–327, 1939.




Fig. 8. Results of the methodology described in Sec. 5.3. On the left, the original object; on the right, the result obtained by our method. First row: A bumped and skewed tube, $\lambda = 20$. Second row: A rocker arm, $\lambda = 8$. Third row: A pelvis bone, $\lambda = 5$.