GRAPHICAL MODELS AND IMAGE PROCESSING
Vol. 57, No. 6, November, pp. 453–461, 1995

Fundamentals of Surface Voxelization

DANIEL COHEN-OR* AND ARIE KAUFMAN

Department of Computer Science, State University of New York at Stony Brook, Stony Brook, New York 11794-4400

Received March 11, 1993; revised March 20, 1995; accepted May 18, 1995

This paper presents the theoretical discrete geometry framework for the voxelization of surfaces. A voxelized object is a 3D discrete representation of a continuous object on a regular grid of voxels. Many important topological properties of a discrete surface cannot be stated solely in terms of connectivity, and thus the concepts of separating, coverage, and tunnel-freeness are introduced. These concepts form the basis for proper voxelization of surfaces. © 1995 Academic Press, Inc.

1. INTRODUCTION

The use of volume rendering technology and voxel-based graphics has recently proliferated, not only for visualizing and analyzing sampled and computed datasets, but also for modeling synthetic scenes [4, 7, 11, 18, 20]. Modeling a geometric scene in voxel space calls for algorithms that generate from a geometric representation of the scene the equivalent discrete voxel-based representation. These algorithms, called *voxelization* (or *3D scan conversion*) algorithms, have received some attention in the literature [2, 8–10, 15]. In this paper, we establish the conditions and the requirements for the proper voxelization of surfaces.

The representation of a continuous object by a discrete set must ensure topological and geometrical consistency. These issues have been successfully dealt with for the approximation of continuous curves in 2D discrete space, mainly because of the linear order a discrete curve inherently possesses [3, 23]. Like a 2D discrete curve, a 3D discrete curve is expressed as a linear sequence of discrete points where the adjacency relation between consecutive points determines the curve connectivity. The connectivity of a curve in 2D space has been proved to be topologically consistent and to satisfy the Jordan theorem [23]. However, there is no straightforward extension to 3D space [17] because connectivity alone fails to describe a 3D discrete surface and there is no natural linear order among the discrete surface points. The topology of 3D discrete spaces has been developed primarily for 3D image processing, where the 3D discrete image is given as a set of voxels [14, 24, 25]. A different approach has been used in the cuberille environment [1, 5], where a surface is composed of a set of voxel faces, representing the boundary between the foreground and the background. In this paper, we adopt the former approach, in which the surface is composed of a set of voxels. We provide conditions on a discrete surface that ensure that its topology is analogous to that of the continuous surface and approximates the continuous surface well, where the main application of these conditions is in the synthesis of voxelized surfaces.

In Sections 2 and 3 we briefly review basic concepts in 3D discrete topology, which aims to mimic the behavior and properties of the corresponding continuous space topology. In Section 4 we propose new criteria for the discrete representation of surfaces, and we examine the relationship among them in Section 5. We develop sufficient conditions for surface separation properties in Section 6. Finally, in Section 7 we discuss other criteria for proper voxelization of surfaces.

2. BASIC DEFINITIONS

Let Z^3 be the subset of the 3D Euclidean space R^3 that consists of all the points whose coordinates are integers. This subset is called the *integer lattice* or the *grid* for short. The *Voronoi neighborhood* of a grid point p is the set of all points in R^3 that are at least as close to p as to any other grid point. The Voronoi neighborhood of a 3D grid point is a closed axis-aligned unit cube known as a *voxel*. The union of all the voxels tessellates R^3, and the interiors of any two voxels are disjoint. Given a function from Z^3 to $\{0, 1\}$, we call the voxels on which the function takes the value 1 "black" voxels, representing objects, and the others are called "white" voxels, representing the transparent background.

Two voxels are 26-*adjacent* if they share just a vertex, just an edge, or just a face. Every voxel has 26 such adjacent voxels defining the 26-neighborhood of the voxel: 8 share a vertex (corner) with the center voxel, 12 share an edge, and 6 share a face. Accordingly, face-sharing voxels are defined as 6-*adjacent*, and edge-sharing and face-sharing

* Current address: School of Mathematical Sciences, Tel-Aviv University, Ramat-Aviv 69978, Israel.

1077-3169/95 $12.00

voxels are defined as 18-*adjacent*. The prefix N is used to define the adjacency relation, where $N \in \{6, 18, 26\}$. An *N-path* in a set of voxels X is a sequence of voxels all in X such that consecutive pairs are N-adjacent. A set of voxels A is *N-connected* if there is an N-path in A between every pair of points in A. An *N-connected component* is a maximal N-connected set.

A common scenario in 2D discrete space is an 8-connected closed curve with an 8-connected "background" penetrating from one side of the curve to the other. To avoid such a scenario, it has been the convention to define opposite types of connectivity for the white and black sets [14]. Opposite types in 2D space are 4 and 8, while in 3D space 6 is "opposite" to 26 or 18. However, for surfaces in 3D, the situation is much more complex because a surface has important topological properties which cannot be stated solely in terms of *N-connectedness*. In this paper, the term "discrete surface" is used for any set of voxels that is intended to be a voxelization of a continuous surface. We focus on and formally define (in Section 4) a *well-voxelized* discrete surface, which is a successful discrete approximation of a continuous surface in terms of its topology and geometry.

3. HOLES, CAVITIES, SIMPLE POINTS, AND PENETRATION

Let Σ be a set of all voxels whose centers form the set

$$\{(x,y,z) \in Z^3 \mid x_0 \le x \le x_1, y_0 \le y \le y_1, z_0 \le z \le z_1\}. \quad (1)$$

The border of Σ is all voxels (x, y, z) such that at least one of the following is true: $x \in \{x_0, x_1\}$, $y \in \{y_0, y_1\}$, $z \in \{z_0, z_1\}$. A set of voxels (pixels) B *N-surrounds* a set of voxels (pixels) A if any N-path $v_0, v_1, \ldots, v_m$, from $v_0 \in A$ to the border of Σ, must meet B (i.e., there is an i such that $v_i \in B$). Let A be a component of $\Sigma - B$ that is N-adjacent to B. When B N-surrounds A in 2D space, we say that A is a *hole* in B. If B N-surrounds A in 3D space, then A is a *cavity* in B.

The 2D analog of a 3D cavity is a hole, but a 3D hole is different from a 3D cavity. A solid torus in 3D has one hole and no cavities, while a hollow torus has an additional hole introduced by the cavity. It is actually difficult to define a hole in 3D; intuitive definitions are given by Kong and Roscoe [12]. If a component has no holes, it is a *simply connected* component. The *Euler number* of a 3D set is defined as the number of components in the set minus the number of holes plus the number of cavities.

A *simple point* is a point whose removal does not change the topology of the set; in particular, it does not change the Euler number of the set. In 2D space, a formal definition is simpler than in 3D since, for example, in 3D the deletion of a voxel might break one hole but might introduce another. For more details, see [16, 26].

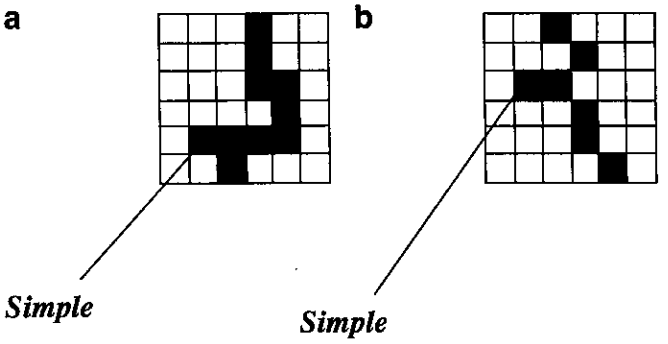

FIG. 1. (a) The black set 8-separates the white set; (b) the black set 4-separates the white set. The two sets contain a simple pixel (voxel). (Whenever the extension to 3D is straightforward, the figures and examples in this paper are shown in 2D rather than in 3D.)

The above definitions attempt to define discrete versions of the topological concepts used in continuous space. The following definitions deal with the penetration of a curve through a surface. Loosely speaking, we use the term continuous surface to denote a set in 3-space that consists of one or more "sheet-like" structures.

A continuous curve that penetrates a continuous surface must meet the surface at some point p. However, a curve that meets a surface does not necessarily penetrate through it; the curve may only meet the surface at a tangent point without passing through it. To distinguish between the two cases, we define $B(p, \varepsilon)$ to be an open ball of radius ε centered at p. A curve α *penetrates* through the surface S at a point p, if for all sufficiently small ε, α meets two different components of $B(p, \varepsilon) - S$. That is, α passes through S at p. However, if α meets neither component, α *coincides* with S at p. In the following sections the term penetration is used for either penetration or coincidence. We have also developed the notion of a discrete penetration, and it will be discussed in Section 4.

Recall that we have defined a voxel as a continuous space unit. Let A be a set of voxels; the union of A, $\cup A$, is a continuous piece of space defined by the discrete set A. The superscript d is added to S to mark some discrete set representing a continuous surface S. It is not an operator nor does S^d have any particular properties. Our goal, however, is to suggest conditions which guarantee S^d to be a discrete analog of the continuous surface S.

4. SEPARATION, COVER, AND TUNNEL-FREENESS

Let S^d be a subset of a set of voxels X. If $X - S^d$ is not N-connected, then S^d is said to be *N-separating* in X (see Fig. 1). We say that a set is *N-separating* if it is N-separating

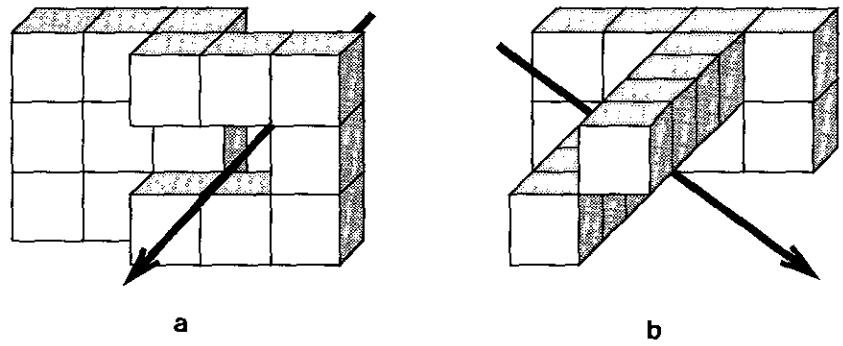

FIG. 2. Two 6-connected sets of voxels. The arrows show where (a) a 6-path and (b) an 18-path of the "background" voxels can pass through the sets.

in Σ. If $X = T \cup Y \cup Z$ and T, Y, and Z are mutually disjoint, and every path connecting a voxel in Y to a voxel in Z meets T, then we say that T N-separates Y and Z in X. The property of N-separating in X is useful when the set of voxels S^d is a voxelization of a continuous surface S such that $\cup X - S$ is not connected (for example, X is equal to Σ and S is an infinite plane passing through $\cup X$). If X is the entire set of voxels Σ, then N-separating is relevant only to closed surfaces or to "infinite" surfaces that partition the space. N-separating does not hold for open objects such as an open cylinder or a disk, or when dealing with surfaces which have boundaries that do not meet the border of Σ.

If S^d is N-separating in Σ, we would then expect S^d to be N-separating in any X such that $S^d \subseteq X$. But some difficulties might arise, especially when X is small and might not meet one of the components of $\Sigma - S^d$. Note that N-separating is a topological property and does not reflect upon how close S^d is to S.

In 2D, an 8-connected simple closed curve is 4-separating and a 4-connected simple closed curve is 8-separating (8-separating is also 4-separating), except in the case of very small curves that are contained in a unit lattice square. This follows from the digital Jordan curve theorem (see Rosenfeld's book [23]). Although 6 is an opposite connectivity type to 26 or 18, there are no analogous results in 3D space [13, 17, 21, 22], where there is no pair (M, N) for which M-connected implies N-separating and there is no natural linear order among the discrete points. A single voxel is a simple example of a 6-connected set that is not 6-separating. Figure 2 presents two applied examples of 6-connected sets, where a background path goes through what intuitively appears as a surface and connects what seems to be the two sides of a 6-connected set. Figure 2a shows a 6-connected set which is not 6-separating (and thus neither 18- nor 26-separating), while Fig. 2b shows a 6-connected set which is not 18-separating (and thus not 26-separating).

Let A be an N-separating set such that $\Sigma - A$ has exactly two N-components. An *N-simple voxel* in A is a voxel v such that $A - v$ is N-separating (see Fig. 1). An N-separating surface is called *N-minimal* if it does not contain any N-simple voxels. A *cover* of a continuous surface is a set of voxels such that every point of the continuous surface lies in some voxel of the cover (see Fig. 3a). S^d covers S if and only if $S \subseteq \cup S^d$. A cover is said to be a *minimal cover* if none of its proper subsets is also a cover. As a consequence, every voxel in a minimal cover meets the continuous surface.

Given a continuous surface S, there is a unique set S^ω called the *supercover* of S, which is the set of all voxels that meet S (see Fig. 3b). Recall that we have defined the extent of a voxel to include its boundary, and the supercover may include voxels that touch the surface at only one point. Thus, the supercover is not necessarily minimal, because the boundary of one voxel meets other voxels.

Let S be a continuous surface such that $R^3 - S$ has exactly two connected components, I and O. Let I^d and O^d be the nonempty set of all voxels wholly contained in I and O, respectively.

THEOREM 1. *S^ω is N-separating in Σ.*

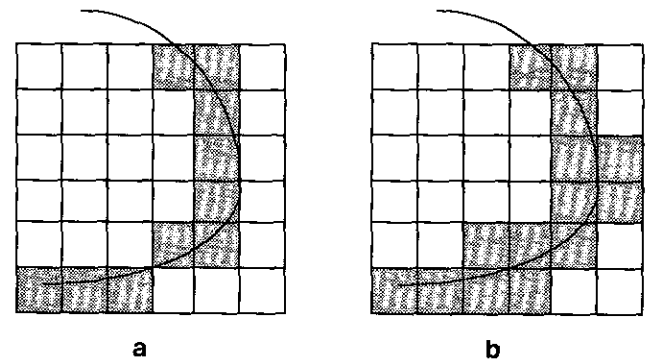

FIG. 3. (a) A cover set of a continuous surface; all its voxels meet the surface, while (b) the supercover of a continuous surface consists of all the voxels that meet the surface.

Proof. The definitions of S^{ω}, I^d and O^d yield $S^{\omega} = \Sigma - (I^d \in O^d)$. It is enough to show that there is no N-path between I^d and O^d, for $N \in \{6, 18, 26\}$. Assume, to the contrary, that there is an N-path Π in $\Sigma - S^{\omega}$ from $v \in I^d$ to $w \in O^d$. Then, there is a continuous curve α wholly contained in $\cup\Pi$ from a point in I to a point in O. Since I and O are two disjoint components of $R^3 - S$, α must meet S at some point $P \in \cup\Pi$. From the definition of S^{ω}, $p \in \cup S^{\omega}$, contradicting the fact that $p \in \cup\Pi$. ■ (Note that the theorem is also true in 2D space.)

Unlike the supercover, a cover, and in particular a minimal cover of S, does not necessarily separate I^d and O^d. There may be some points on S which meet more than one voxel (e.g., Fig. 3b) and thus can be covered by one voxel, but at the same time can be included in other voxels. These singular points are defined as follows: a point p is said to be a *singular point* in $\cup S^d$ if $\cup V(p, S^d) - \{p\}$ is not simply connected, where $V(p, S^d)$ is the subset of S^d of all voxels that include p.

Let S^d be the set of voxels generated by some voxelization algorithm of S. The voxelization of S strives to include only voxels close to S, that is, $S^d \subseteq S^{\omega}$. Thus, we define a *partial cover* as a connected subset of a minimal cover. If S^d is a partial cover of S, then $v \in S^d \Rightarrow v \cap S \neq \emptyset$. Consequently, we require S^d to be a partial cover of S to prevent the voxelization from being too coarse. Note that a partial cover is not unique. While we do not want to impose many restrictions of S^d, we would clearly like S^d to N-separate I^d and O^d.

In our discussion, the meaning of I and O is the inside and outside of a surface. As mentioned before, N-separating fails to deal with borders, in which case the surface does not have two sides. The following definition imposes a separation property without being restricted to surfaces that have two sides. In essence, it prevents a discrete penetration of the background voxels where there is no analogous continuous penetration.

First, let $C(\Pi)$ be the union of straight line segments that join the centers of consecutive voxels in an N-path Π. Thus $C(\Pi)$ is a polygonal arc (a polyline) of $k - 1$ elements, where k is the number of voxels in the N-path Π. A set of voxels S^d is *N-tunnel-free with respect to S* if for every N-path Π in $\Sigma - S^d$, $C(\Pi)$ does not penetrate S. By an *N-tunnel* through S^d, we mean an N-path Π in $\Sigma - S^d$ such that $C(\Pi)$ penetrates S.

If in the above definition of N-tunnel-free we replace "$C(\Pi)$ does not penetrate S" with "there is no (continuous) curve in $\cup\Pi$ that penetrates S," then a set of voxels that is N-tunnel-free with respect to S would necessarily be a cover of S. An advantage of the above definition of N-tunnel-free is that it allows us to apply the concept to voxelizations which are not covers. A coverage requirement is especially restrictive for thin surfaces such as 6-tunnel-free surfaces (see 2D examples in Fig. 4). The continuous curve $C(\Pi)$ enables penetration through some limited parts of S. However, there are some situations where $C(\Pi)$ penetrates very small parts of S which are close to the surface edge (see example in Fig. 5).

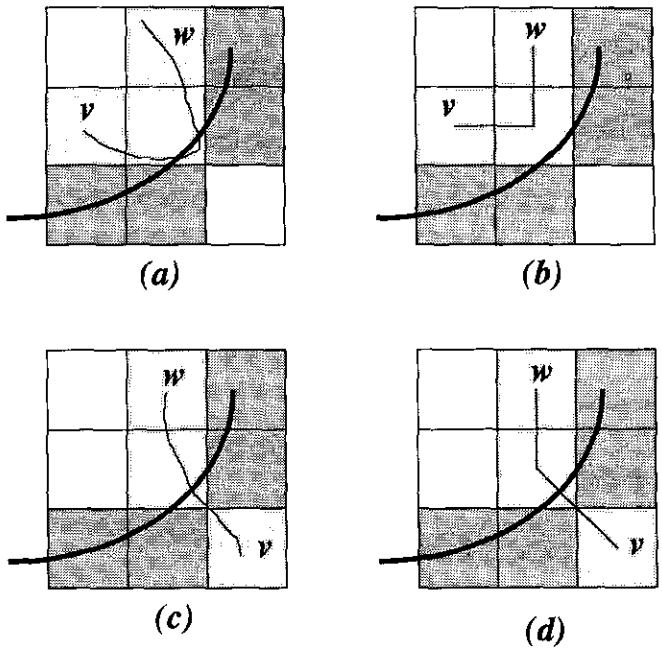

FIG. 4. A path (in light gray) from w to v. An unrestricted curve in the tunnel may penetrate the continuous surface but not necessarily the discrete surface (a), but the polygonal arc avoids small nonsignificant penetrations (b). In the case of a true penetration (c), the polygonal arc also penetrates the continuous surface (d).

Note that the voxelization of an N-tunnel-free S^d is not only with respect to its topology (separating) but also with respect to geometry, as it guarantees the discrete surface to be close to (almost covering) the continuous surface, which is a geometric property. Actually, only small protru-

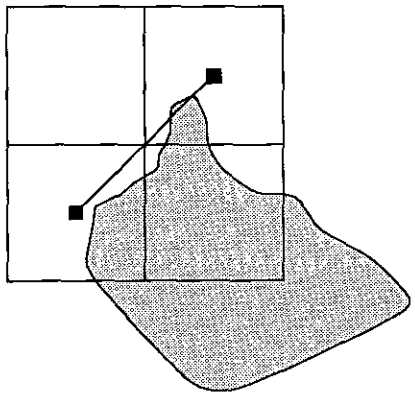

FIG. 5. The continuous curve $C(\Pi)$ might penetrate small parts of the continuous surface which are close to the surface edge.

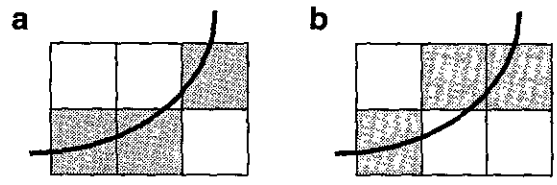

FIG. 6. Two possible discrete representations of the same surface. (a) seems to better approximate the surface than (b).

sions within a voxel (Fig. 5) are allowed by the tunnel-free definition. A discrete surface S^d, which is an N-tunnel-free partial cover of a connected surface S, is said to be *well voxelized.*

A well-voxelized surface is not unique; for example, the two voxelizations of the curve in Fig. 6 are tunnel-free (they are 4-tunnel-free in 2D space). Intuitively, that in Fig. 6a seems to better approximate the curve, while, formally, a metric needs to be defined. One possible metric is $Q = \Sigma_{v \in S^d} D(v, S)$, where $D(v, S)$ is the (minimal) Euclidean distance between the center of v and the surface S. Then, the best voxelization of S minimizes Q over all possible well-voxelizations of S.

Unlike the definition of separating, the definition of tunnel-free is valid for any continuous surface and does not assume that the continuous surface has two sides. Clearly, the existence of 6-tunnels implies the existence of 18-tunnels, and the existence of 18-tunnels implies the existence of 26-tunnels. Similarly, the absence of 26-tunnels implies the absence of 18-tunnels, which implies the absence of 6-tunnels. Whenever a surface is 26-tunnel-free, we simply write tunnel-free with no prefix since all types of tunnels are absent. Now we can state theorems stronger than Theorem 1:

THEOREM 2. *S^ω is tunnel-free.*

Proof. Assume, to the contrary, that there is an N-path Π in the complement of S^ω which contains a continuous path which meets S at a point p. However, from the definition of S^ω we get $p \in \cup S^\omega$, contradicting the fact that $p \in \cup \Pi$. ■

COROLLARY. *Let* w *be a set of voxels such that* $S \cap \cup(\Sigma - w) = \emptyset$. w *is tunnel-free.*

Proof. $S^\omega \subseteq w$, and the property of being tunnel-free is preserved under the superset operation. ■

THEOREM 3. *Let* S *be a continuous surface such that* $\cup\Sigma - S$ *has two components,* I *and* O. *We assume that there are two voxels,* w *and* v, *that do not meet the surface and* $w \subseteq I$ *and* $v \subseteq O$. *Let* S^d *be* N*-tunnel-free with respect to* S. S^d *is* N*-separating.*

Proof. We have to show that $\Sigma - S^d$ has two components. It is enough to show that w and v are not connected by any N-path in $\Sigma - S^d$. Assume to the contrary that there is an N-path Π connecting w and v in $\Sigma - S^d$. Then the continuous curve $C(\Pi)$ connects w and v and must penetrate S, contradicting the definition of N-tunnel-free. ■

We have mentioned before that there are many important topological properties that cannot be stated only in terms of N-connectedness. In 2D, we have the two following results which describe the relation between connectivity and tunnel-freeness:

Let Π be a set of pixels such that every pair of pixels in Π that are 8-adjacent but not 4-adjacent has a common 4-neighbor in Π. Then Π is 8-tunnel-free with respect to any set $\alpha \subseteq \cup\Pi$.

Let Π be a set of pixels and let α be a plane set such that each point in α is less than 1 unit away from the center of a pixel in Π. Then Π is 4-tunnel-free with respect to α.

There are no analogous results for 3D surfaces, but we can get a weaker result. Let S^d be a 6-connected set of voxels, and let $S \subseteq \cup S^d$ be some continuous surface contained in $\cup S^d$. Regard S^d as some voxelization of S; then S^d is 6-tunnel-free. This holds because 6-adjacency does not have singular points, as 8-, 18-, and 26-adjacency have.

5. THE RELATIONSHIP BETWEEN DIFFERENT SETS OF VOXELIZED SURFACES

We have introduced above the separating and coverage properties and defined tunnel-free surfaces. We have also defined the supercover and proved that it is separating and tunnel-free. Figure 7 depicts the relationships among these discrete surface sets (for the voxelizations of typical closed surfaces). Separation implies neither coverage nor tunnel-freeness. Separation relates only to the topology preservation of the discrete surface with respect to the continuous one. On the other hand, as mentioned above, a cover does not guarantee separation. Figure 8 shows the relationships among the tunnel-free sets. Note that coverage implies 6-tunnel-freeness.

For a particular continuous surface S, such that $R^3 - S$ has exactly two connected components, I and O, there are

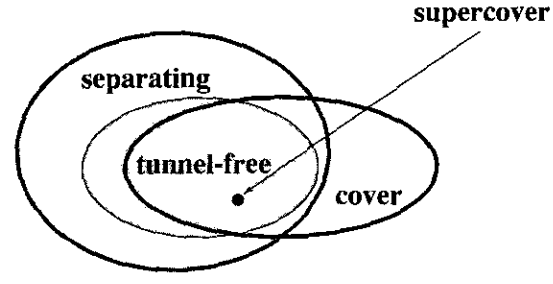

FIG. 7. The relationship among different voxelizations of closed surfaces.

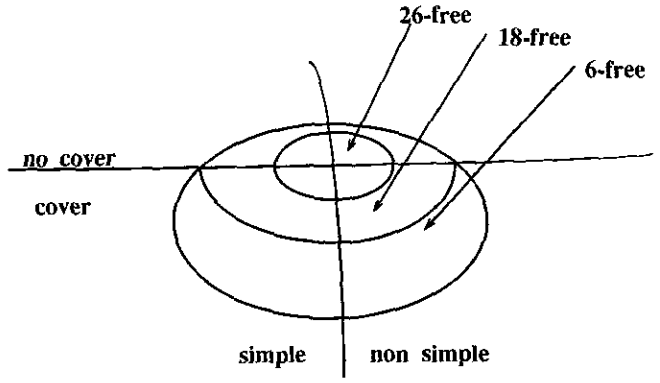

FIG. 8. The relationship among different sets of tunnel-free surfaces.

many sets that satisfy the coverage property and many that satisfy the separating property. Let S^d be an N-separating minimal cover of S. In most cases, S^d is also N-tunnel-free, but since the property of separating is not local and tunnel-freeness may not preserve the original topology of S, S^d is not necessarily N-tunnel-free.

A tunnel-free surface is not necessarily a cover because a small (uncovered) protrusion of the continuous surface into a voxel is not considered by the tunnel-free definition (as in Fig. 9). However, except for small protrusions, 26- and 18-tunnel-free surfaces cover the continuous surface. But a 6-tunnel-free surface, similar to an 8-connected path, skips over small segments of the continuous object and does not provide a cover.

A tunnel-free surface is not necessarily separating because the original continuous surface may not be closed. A tunnel-free voxelized surface provides local separation (topology preserving) and is geometrically close enough to the continuous surface. If coverage is desired too, a more restrictive voxelization is needed and, consequently, more computation is involved. Assume a surface S such that $\Sigma - S$ has exactly two components. In practice an efficient voxelization of S, in terms of time and space, should produce an N-minimal surface. An N-minimal N-tunnel-free surface, with or without the coverage property, is a successful voxelization, but this is more restricted than a well-voxelized surface, which is just a tunnel-free partial cover of S.

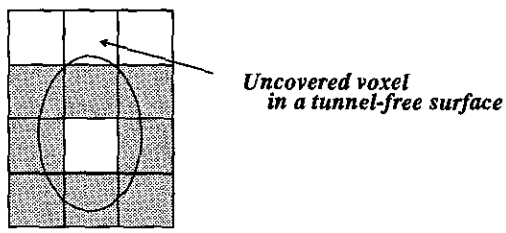

FIG. 9. A small part of a tunnel-free surface may not be covered.

6. SLICEWISE SEPARATING PROPERTIES

Ideally, the voxelization S^d of a continuous surface S is an N-tunnel-free partial cover. However, such a voxelization algorithm can be computationally expensive. Separating, on the other hand, is easier to achieve. The exact location of the continuous surface is not required for separation. Separation is defined in order to separate two sets, where the actual goal is to separate the two sides of the continuous surface. However, it is not easy to know what the two sides of a surface are just by a local operation. Moreover, some surfaces, such as a Mobius strip, have only one side. In this section, new sufficient conditions for N-separating are developed based on 2D properties, which are easier to provide. A *slice* is defined as a set of all voxels with one constant coordinate component. The slice is, thus, parallel to one of the primary axis planes. We denote by C^x an x-slice parallel to the yz primary plane. Similarly, a slice C^y and a slice C^z are parallel to the xz and xy planes, respectively.

THEOREM 4. *Let* T^d *be a set of voxels and let* I^d *and* O^d *be a partition of* $\Sigma - T^d$. *If for every* 1*-slice* C^l *such that* $l \in \{x, y, z\}$, $T^d \cap C^l$ 4*-separates* $I^d \cap C^l$ *and* $O^d \cap C^l$, *then* T^d 6*-separates* I^d *and* O^d.

Proof. Assume to the contrary that there is a 6-path connecting I^d and O^d. Then there must be two 6-adjacent voxels v and w in the 6-path, such that $v \in I^d$ and $w \in O^d$. v and w differ by only one coordinate value, and thus, there is a slice C^{l_i} such that $v \cup w \in C^{l_i}$, contradicting the fact that $T^d \cap C^l$ 4-separates $I^d \cap C^l$ and $O^d \cap C^l$, for any slice C^l. ■

The requirements in Theorem 4 can be relaxed by considering only two types of slices, because if v and w are 6-adjacent, then $v \cup w \in C^{l_1}$ and $v \cup w \in C^{l_2}$, where $l_1 \neq l_2$. That is, the 6-path can be detected along either l_1 or l_2. By testing only two (out of three) types of slices, it is guaranteed that any 6-path can be detected. Without loss of generality, the Z-slices can be avoided.

THEOREM 5. *Let* T^d *be a set of voxels and let* I^d *and* O^d *be a partition of* $\Sigma - T^d$. *If for every* 1*-slice* C^l, *such that* $l \in \{x, y\}$, $T^d \cap C^l$ 4*-separates* $I^d \cap C^l$ *and* $O^d \cap C^l$, *then* T^d 6*-separates* I^d *and* O^d.

Theorem 4 can be extended to characterize 18-separating:

THEOREM 6. *Let* T^d *be a set of voxels and let* I^d *and* O^d *be a partition of* $\Sigma - T^d$. *If for every* 1*-slice* C^l *such that* $l \in \{x, y, z\}$, $T^d \cap C^l$ 8*-separates* $I^d \cap C^l$ *and* $O^d \cap C^l$, *then* T^d 18*-separates* I^d *and* O^d.

Unfortunately, 26-separating is inherently a 3D property and a 26-path cannot be characterized by inspecting slices.

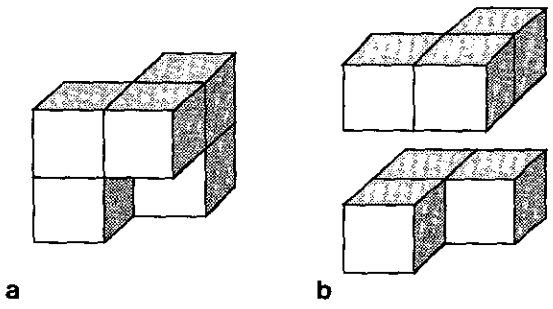

FIG. 10. (a) A 6-connected set with a 26-tunnel, (b) consists of two "slices" of 4-connected paths.

The $2 \times 2 \times 2$ configuration of voxels shown in Fig. 10 is an example of a 6-connected object whose intersection with any slice is a 4-connected path; yet there is a 26-connected path of background voxels through the discrete surface. This special configuration is denoted by H^{26}. This H^{26} is the only $2 \times 2 \times 2$ configuration (up to rotations and mirror transformations) which does not contain an 18-path but does contain a 26-path:

THEOREM 7. *Let* T^{d} *be an* 18*-separating set of voxels. If none of its* $2 \times 2 \times 2$ *subsets form an* H^{26} *in any orientation, then* T^{d} *is* 26*-separating.*

Proof. Immediately from the definition of H^{26}. ∎

We can also define diagonal slices by which 26-separating can be characterized. Let a (*minor*) *diagonal slice* γ be the set of voxels (x, y, z) such that the difference between two of its coordinates is constant (e.g., $x - y = \gamma$). The set of voxels of a β slice can be mapped onto a 2D slice C^F by a one-to-one function F. For example, a diagonal slice $x - z = \gamma$ can be "flattened" onto a 2D slice $F(x, y, z) = (x, y)$. Three-dimensional corner adjacent voxels in γ are 2D-corner adjacent in C^F. Thus, a 26-tunnel is mapped onto an 8-tunnel in C^F:

THEOREM 8. *Let* T^{d} *be a set of voxels and let* I^{d} *and* O^{d} *be a partition of* $\Sigma - T^{\mathrm{d}}$. *Let* T^{d} 18*-separate* I^{d} *and* O^{d}. *If for every diagonal slice* γ, $\mathrm{F}(\mathrm{T}^{\mathrm{d}} \cap \gamma)$ 8*-separates* $\mathrm{F}(\mathrm{I}^{\mathrm{d}} \cap \gamma)$ *and* $\mathrm{F}(\mathrm{O}^{\mathrm{d}} \cap \gamma)$, *then* T^{d} 26*-separates* I^{d} *and* O^{d}.

7. TOWARD WELL VOXELIZATION

In this section we develop sufficient criteria for the voxelization of an N-tunnel-free surface, aiming at generating a minimal tunnel-free coverage, which can be referred to as a well voxelization of the continuous surface S. The voxelized set S^{d} can be exactly the supercover and thus is guaranteed to be tunnel-free. However, for some surfaces S, such voxelization is too coarse because some points on the continuous surface are met by up to eight voxels, and the voxelization of S might be too thick with an excess of simple voxels. That is, there is a subset of the supercover which is still tunnel-free. For the purpose of deriving such a subset, the voxelization algorithm assumes a refined definition of a *reduced-voxel* marked *R-voxel*. The new refined definition reduces the number of voxels that can meet the same point. Recall that the voxel has been defined as a closed unit cube that includes its boundary. The aggregate of all voxels tessellates the space where every face of a voxel meets and is shared by the face of a neighboring voxel. Similarly, the edge of one voxel is common to three other voxels, and a corner of one voxel meets seven other voxels. The idea is to reduce the amount of overlapping among adjacent voxels in a way that preserves the tunnel-free property of the supercover, while on the other hand, ensuring that each voxel in the new supercover is not an N-simple voxel.

Let an R-pixel be the unit square area centered at a lattice point which includes only a part of its boundary (as shown in Fig. 11): two of the edges, one horizontal and one vertical, and only two of its corners, one from each diagonal. An R-pixel at integer coordinates (i, j) is denoted by $P(i, j)$, which is the set

$$\{(x, y) \mid i - 0.5 < x \leq i + 0.5, j - 0.5 < y \leq j + 0.5\} - \{(i + 0.5, j - 0.5)\}. \quad (2)$$

Now each edge meets only one R-pixel and each corner point meets exactly two R-pixels. Consequently, with the R-pixel boundary definition, the new supercover, denoted by $S^{\hat{\omega}}$, includes fewer pixels, but Theorem 2 still holds because $S \subseteq S^{\hat{\omega}}$, and $S^{\hat{\omega}}$ is tunnel-free as the following theorem states.

THEOREM 9. *Given a continuous set* $S \in R^2$, *let* $S^{\hat{\omega}}$ *be the pixelization of* S *such that* $S^{\hat{\omega}} = \{P(i, j) \mid P(i, j) \cap S \neq \emptyset\}$. $S^{\hat{\omega}}$ *is tunnel-free.*

Proof. Assume to the contrary that there is an N-tunnel Π in the complement of $S^{\hat{\omega}}$ such that $C(\Pi)$ meets S at a point p. From Theorem 2 we know that S^{ω} is tunnel-free.

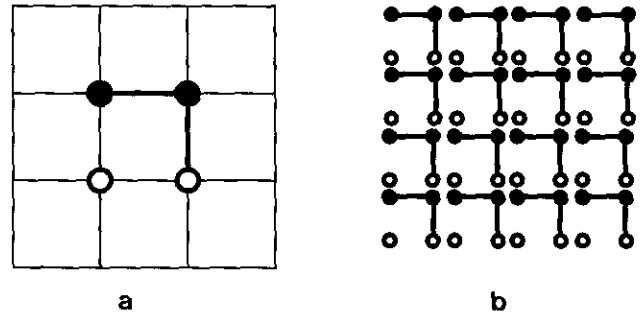

FIG. 11. (a) The boundary of a single pixel, and (b) the relationship among adjacent pixels.

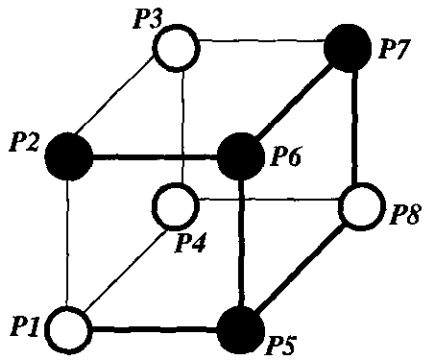

FIG. 12. A possible definition for the R-voxel boundary: all faces incident with P6, the thick edges (P6, P2), (P6, P7), (P6, P5), (P5, P8), (P5, P1), and (P8, P7), and all black vertices (P6, P7, P5, P2).

Thus, p must be in $S^{\omega} - S^{\hat{\omega}}$, which is on the exterior of a pixel. Therefore, p is on some segment connecting two consecutive pixels v and w in $C(\Pi)$. v and w are either edge adjacent or corner adjacent. If v and w are edge adjacent, then p is on the edge between them. Since by definition an R-pixel contains one vertical edge and one horizontal edge, one of the reduced pixels v or w must meet S at p and must be included in $S^{\hat{\omega}}$, contradicting the assumption. If v and w are corner adjacent, then p is on the pixel corner connecting between them. Since by definition an R-pixel contains one corner on each main diagonal, one of the reduced pixels must meet S at p and must be included in $S^{\hat{\omega}}$, contradicting the assumption. ■

Note that two R-pixels still overlap at an upper corner point. Taking only one corner prevents tunnels if the same connectivity is used for the background pixels and the foreground pixels [19], but we do not want to limit the behavior of the background. The new definition of the R-pixel boundary is aimed only at redefining the set S^{ω} to the set $S^{\hat{\omega}}$. In order to make $S^{\hat{\omega}}$ a canonical rasterization [3] (loosely speaking, the canonical rasterization of a 2D curve is the minimal set of 4-connected pixels pierced by the continuous curve), a "patch" has to be added to the $S^{\hat{\omega}}$ definition. A similar approach has been used by Hobby [6], who has added pixels in the special case of ambiguity. Rather than adding pixels when necessary, we can omit pixels when unnecessary. The overlapping can be regarded as caused by the upper left corner. This closed corner has been added to prevent tunnels when the continuous line passes between diagonally adjacent R-pixels.

We define the R-voxel boundary in 3D in a similar way. The unit cube boundary contains three pairs of faces, six pairs of edges, and four pairs of corner points. One of each pair is selected to be included in the new R-voxel definition. A possible definition is given in Fig. 12. With this definition, the overlapping between adjacent R-voxels is reduced and the superset is thus much thinner but still 26-tunnel-free. Thus, we mark R with the subscript 26 to indicate that with the R_{26}-voxel the superset is 26-tunnel-free. However, in rare cases, it may contain simple voxels.

Some modification to the R-voxel definition can yield even thinner surfaces, which guarantees only 18-tunnel-free or just 6-tunnel-free. We can define an R_{18}-voxel similarly to an R_{26}-voxel, except that all the corner points are excluded from the R_{18}-voxel boundary. Then every point in R^3 meets one R_{18}-voxel either inside or on a face, meets two R_{18}-voxels on an edge, or is located at the corner of eight adjacent R_{18}-voxels, where it is actually in none of them. Assume a continuous surface S that includes more than one 3D point. If $S^{\hat{\omega}}$ is not tunnel-free, some 18-path Π must contain a continuous path $C(\Pi)$ which meets S, say at p. Since $p \in \cup\ (\Sigma - S^{\hat{\omega}})$ and $p \in S$, from the definition of the R_{18}-voxel boundary and the definition of $S^{\hat{\omega}}$, p must be on a corner of eight voxels. However, this results in a contradiction since Π is an 18-path, and the point p must be on a face, on an edge, or in the inner volume of a voxel.

Excluding all edges from the R_{18}-voxel boundary, we get the R_6-voxel which, in a similar way, makes $S^{\hat{\omega}}$ a 6-tunnel-free surface. This is, however, sometimes too restrictive: if only 6-connectivity is used for the background, then $S^{\bar{d}}$ can be a thinner set and not necessarily $S \subseteq \cup\ S^{\bar{d}}$.

In summary, we have found sufficient conditions for a set of voxels to be N-tunnel-free by creating a cover while reducing the number of N-simple voxels. For 18- and 26-tunnel-free surfaces, the N-simple voxels can be avoided with some local handling, as mentioned above. However, other definitions should be developed to generate a minimal or even a close-to-minimal 6-tunnel-free surface.

8. CONCLUSIONS

We have seen that a discrete voxelized surface cannot be properly described by its connectivity alone because this does not imply a proper voxelization of a surface. Separation, however, is just a topological property. We have thus introduced the coverage property, which indicates that the voxelized surface is "close" to the continuous surface. Since coverage does not guarantee separation and also might be too restrictive, we have defined tunnel-free surfaces which have stronger properties. In fact, tunnel-freeness is in some sense a local analog of the global property of "separating." Tunnel-freeness also guarantees that the discrete surface is close to the continuous surface without necessarily being a coverage. We have also shown that slicewise properties provide necessary conditions for separation, and we have devised conditions for a simple cover, which guarantee tunnel-free surfaces. The concepts developed in this paper form a discrete geometry framework for devising algorithms for properly voxelizing surfaces.

ACKNOWLEDGMENTS

This work has been partially supported by the National Science Foundation under Grant CCR-9205047. We are grateful to T. Yung Kong for numerous discussions and his invaluable insight and help in converting some intuitions into correct mathematical form. Thanks also to Dan Gordon and Ari Rappaport for their useful comments. Special thanks to the anonymous referees for their invaluable advice.

REFERENCES

1. E. Artzy, G. Frieder, and G. T. Herman, The theory, design, implementation and evaluation of a three dimensional surface detection algorithm, *Comput. Graphics Image Process.* **15,** 1981, 1–24.
2. D. Cohen and A. Kaufman, Scan conversion algorithms for linear and quadratic objects, in *Volume Visualization* (A. Kaufman, Ed.), pp. 280–301, IEEE Computer Society Press, Los Alamitos, CA, 1991.
3. H. Freeman, Computer processing of line-drawing images, *ACM Comput. Surveys* **6**(1), 1974, 57–97.
4. N. Greene, Voxel space automata: Modeling with stochastic growth processes in voxel space, *Comput. Graphics* **23**(3), 1989, 175–184.
5. G. T. Herman and D. Webster, A topological proof of a surface tracking algorithm, *Comput. Vision Graphics Image Process.* **23,** 1983, 162–177.
6. J. D. Hobby, Rasterization of nonparametric curves, *IEEE Comput. Graphics Appl.* **9**(3), 1990, 262–277.
7. J. T. Kajiya and T. L. Kay, Rendering fur with three dimensional textures, *Comput. Graphics* **23**(3), 1989, 271–280.
8. A. Kaufman and E. Shimony, 3D scan-conversion algorithms for voxel-based graphics, *Proceedings, ACM Workshop on Interactive 3D Graphics, Chapel Hill, NC, October 1986*, pp. 45–76.
9. A. Kaufman, Efficient algorithms for 3D scan-conversion of parametric curves, surfaces, and volumes, *Comput. Graphics* **21**(4), 1987, 171–179.
10. A. Kaufman, An algorithm for 3D scan-conversion of polygons, in *Proceedings, EUROGRAPHICS' 87, Amsterdam, Netherlands, August 1987*, pp. 197–208.
11. A. Kaufman, D. Cohen, and R. Yagel, Volume graphics, *IEEE Computer* **26**(7), 1993, 51–64.
12. T. Y. Kong and A. W. Roscoe, Characterizations of simply-connected finite polyhedra in 3-space, *Bull. London Math. Soc.* 1985, 575–578.
13. T. Y. Kong and A. W. Roscoe, A continuous analog of axiomatized digital surfaces, *Comput. Vision Graphics Image Process.* **29,** 1985, 60–86.
14. T. Y. Kong and A. Rosenfeld, Digital topology: Introduction and survey, *Comput. Vision Graphics Image Process.* **48,** 1989, 357–393.
15. W. Mokrzycki, Algorithms of discretization of algebraic spatial curves on homogeneous cubical grids, *Comput. & Graphics* **12**(3/4), 1988, 477–487.
16. D. G. Morgenthaler, Three-dimensional simple points: Serial erosion, parallel thining, and skeletonization, Technical Report-1005, Computer Science Center, University of Maryland, College Park, 1981.
17. D. G. Morgenthaler and A. Rosenfeld, Surface in three-dimensional digital images, *Inform. Control* 1981, 227–247.
18. V. A. Norton, Generation and rendering of geometric fractals in 3-D, *Comput. Graphics* **16**(3), 1982, 61–67.
19. T. Pavlidis, *Algorithms for Graphics and Image Processing*, Computer Press, Rockville MD, 1982.
20. K. Perlin and E. M. Hoffert, Hypertexture, *Comput. Graphics* **23**(3), 1989, 253–262.
21. G. M. Reed and A. Rosenfeld, Recognition of surfaces in three-dimensional digital images, *Inform. Control* 1982, 108–120.
22. G. M. Reed, On the characterization of simple closed surfaces in three-dimensional digital images, *Comput. Vision Graphics Image Process.* **25,** 1984, 226–235.
23. A. Rosenfeld, Digital geometry, in *Picture Languages*, Academic Press, New York, 1979.
24. A. Rosenfeld, Three-dimensional digital topology, *Inform. Control* 1981, 119–127.
25. J. Toriwaki, S. Yokoi, T. Yonekura, and T. Fukumura, Topological properties and topology-preserving transformation of a three-dimensional binary picture, in *Proceedings, 6th International Conference on Pattern Recognition*, IEEE82CH1801-0, 1982, pp. 414–419.
26. Y. F. Tsao and K. S. Fu, A 3D parallel skeletonwise thining algorithm, in *Proceedings, IEEE PRIP Conference*, 1982, pp. 678–683.