Computers & Graphics 46 (2015) 36–47

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

SMI 2014

Chained segment offsetting for ray-based solid representations

Jonas Martinez *, Samuel Hornus, Frédéric Claux, Sylvain Lefebvre

Inria, Loria, France

ARTICLE INFO

Article history:
Received 7 July 2014
Received in revised form
22 August 2014
Accepted 17 September 2014
Available online 2 October 2014

Keywords:
Morphological operations
3D modeling
Fabrication

ABSTRACT

We present a novel approach to offset solids in the context of fabrication. Our input solids can be given under any representation: boundary meshes, voxels, indicator functions or CSG expressions. The result is a ray-based representation of the offset solid *directly used for visualization and fabrication*: we never need to recover a boundary mesh in our context.

We define the offset solid as a sequence of morphological operations along line segments. This is equivalent to offsetting the surface by a solid defined as a Minkowski sum of segments, also known as a *zonotope*. A zonotope may be used to approximate the Euclidean ball with precise error bounds.

We propose two complementary implementations. The first is dedicated to solids represented by boundary meshes. It performs offsetting by modifying the mesh in sequence. The result is a mesh improper for direct display, but that can be resolved into the correct offset solid through a ray representation. The major advantage of this first approach is that no loss of information – re-sampling – occurs during the offsetting sequence. However, it applies only to boundary meshes and cannot mix sequences of dilations and erosions. Our second implementation is more general as it applies directly to a ray-based representation of any solid and supports any sequence of erosion and dilation along segments. We discuss its fast implementation on modern graphics hardware. Together, the two approaches result in a versatile tool box for the efficient offsetting of solids in the context of fabrication.

1. Introduction

Morphological operations [1] – such as erosions and dilations – are important operations in solid modeling [2,3]. In the context of fabrication, erosions and Boolean differences can be used for example to hollow a solid or create a mold, while closing operations can remove small holes in a model. Fig. 1 illustrates a few morphological operations obtained by our method.

Many approaches consider the *offset surface* obtained after the dilation or erosion of the solid by the Euclidean ball of radius d centered at the origin. The offset surface is the set of points at distance d from the object boundary. The exterior (resp. interior) offset is the subset of the offset surface lying outside (resp. inside) the solid. The exact computation of offset surfaces for general inputs is difficult. Therefore, a number of approximations have been proposed (see Section 2). However, many of these approximations either restrict the type of input, perform aggressive re-sampling, or require computationally heavy and relatively complex algorithms [4].

In this work we consider sequences of erosions and dilations along line segments. It is worth noting that the result of a sequence of *dilations* along segments is equivalent to a Minkowski sum between the solid and an object known as a *zonotope*. The zonotope is defined as the Minkowski sum of the set of segments.

A zonotope is usually sufficient for our target applications in manufacturing: the main differences with a ball are essentially aesthetic (see Fig. 12), and often only impact hidden surfaces when used for molds and hollowing. Nevertheless, there are known algorithms to approximate a ball with a zonotope within a prescribed error bound [5,6]. Sequences of erosions and dilations along line segments therefore provide a general framework to perform complex morphological operations. This includes closings and openings, obtained by mixing dilations and erosions in sequence.

Our work is focused on obtaining ray-based solid representations [7] for direct visualization and fabrication – typically through slicing and additive manufacturing [8,9]. We do not attempt to recover a boundary representation of the result. Our modeler takes any solid representation as an input – boundary meshes, voxels, CSG expressions – and converts them into ray-based representations for visualization and fabrication. The conversion occurs at the resolution of the screen or manufacturing process, therefore minimizing the loss of information due to sampling. Our modeler

* Corresponding author.
E-mail address: jonas.martinez-bayona@inria.fr (J. Martinez).

http://dx.doi.org/10.1016/j.cag.2014.09.017
0097-8493/

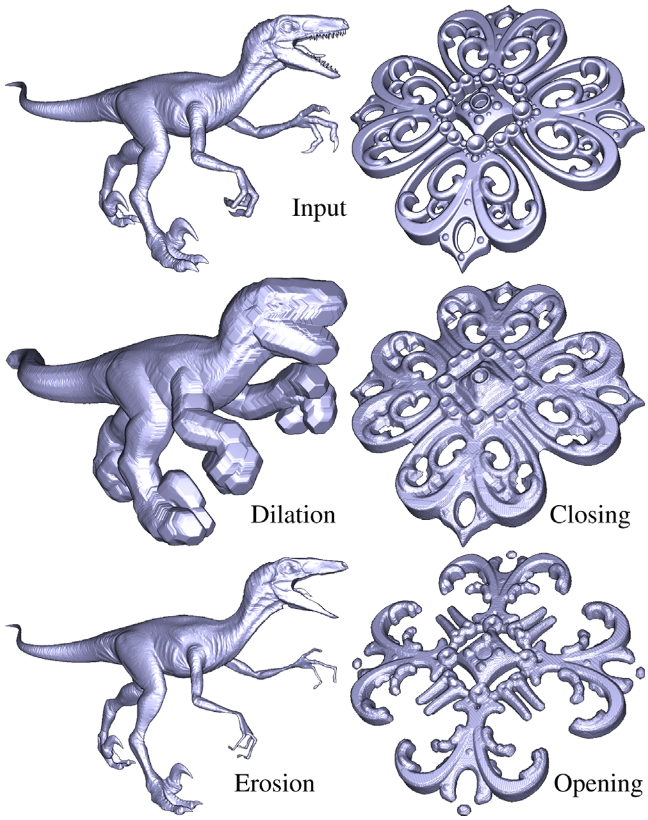

Fig. 1. Segment morphological operations with the specialized approach for meshes (left column) and the generic algorithm (right column). The mesh approach shows the dilation and erosion with a truncated octahedron (zonohedra, see Section 6). Notice the small erosion successfully applied to the raptor model. The generic approach is used to perform a closing and an opening also with the truncated octahedron.

is based on a fast GPU implementation, enabling the construction of ray-based representations at high resolutions (see Section 7).

Contributions: The key observation of our work is that the ray representation of solids is amenable to a simple and fast implementation of morphological operations with *line segments*, and as a consequence, to sequences of morphological operations with zonotopes. To the best of our knowledge, no previous work considers morphological operations between zonotopes and ray-based solid representations. Unlike most of the existing methods, our technique avoids any explicit treatment of topological changes. Erosions and dilations can be combined in any order to achieve complex operations.

We propose two complementary techniques. First, in Section 4 we introduce an efficient algorithm to perform morphological operations on a ray-based representation of a solid. The advantage of this approach is that it applies to any solid that can be captured by a ray-based representation. Its drawback stems from the sampling resolution that approximates the solid at each step. We discuss error bounds for the process in Section 4.3. In our context, and thanks to the high computational efficiency of the presented technique, we can afford the use of a resolution matching that of the manufacturing process of the final object.

Second, we propose in Section 5 a specialized approach for boundary meshes, which postpones the conversion to a ray-based representation to *after* an entire sequence of dilations or sequence of erosions, thereby removing any re-sampling error due to intermediate steps.

The time complexity of the presented algorithms is bounded by the complexity of the solid surface, instead of its volume. Thus, their performance is expected to scale better than voxelization methods. We provide an implementation of all of our algorithms which are both simple to implement and highly parallel.

2. Related work

This section reviews existing approaches for the computation of offset surfaces in general, and then focuses on methods using ray-based representations.

Computing offset surfaces: Early approaches rely on convolutions to compute offset surfaces and Minkowski sums. These methods obtain a superset of primitives of the offset surface that are trimmed and filtered to form the final boundary [10]. Evans and Koppelman [11] compute the Minkowski sum of a polyhedral object along a sequence of translational sweeps, and propose to approximate the Euclidean ball with a zonotope for surface offsetting. To the best of our knowledge this is the only previous approach that considers zonotopes for morphological operations, but it focuses on generating polyhedral results while our focus is on ray-representations. Kaul and Rossignac [12] presented a set of criteria to filter the primitives that do not belong to the Minkowski sum. Peternell and Steiner [13] presented a convolution algorithm for objects with piecewise boundaries. Campen and Kobbelt [14] introduced an exact approach for Minkowski sums between polyhedra that also culls a superset of primitives. Convolution methods usually suffer from geometric robustness issues.

The offset surface can also be extracted from the distance field of the object surface, as it implicitly represents offset surfaces. Frisken et al. [15] presented the adaptively sampled distance-fields, which among other operations, is able to perform surface offsetting. Varadhan and Manocha [16] approximate the Minkowski sum with a distance field isosurface extraction, guaranteeing a Hausdorff distance bound on the approximation. Pavić and Kobbelt [17] traverse an octree and split each cell which is potentially intersected by the offset surface, in order to recover it. Lee et al. [18] presented an accurate method to compute the distance field, which is able to render offset surfaces by considering a union of balls. The main drawback of distance field methods is that they usually require high amount of memory in order to ensure accuracy.

Offset surfaces can also be computed from point-based representations. Chen et al. [19] generate a set of candidate points that are used to obtain a voxelization of the offset surface. Lien et al. [20] and Netaluri and Shapiro [21] perform the Minkowski sum between two point-based surfaces. The approach explicitly distinguishes the interior and boundary points of the Minkowski sum. Recently, Calderon and Boubekeur [22] introduced a morphological analysis framework for point clouds, which is able to perform morphological dilations and erosions. These operations remain expensive on point sets as the interior of the solid is not explicitly available.

A last family of methods generates a voxelization of the offset surface. Li and McMains [23,24] and Leung et al. [25] presented GPU approaches to compute the Minkowski sum of polyhedra by computing pairwise Minkowski sums, and obtaining a voxelization of its union. The memory requirements of these methods rise rapidly as the voxelization resolution increases. In addition, the error tends to be larger than that of a ray-based approach where the sampling directions can freely vary.

Surface offsetting with ray-based representations: To the best of our knowledge, the *dexel structure* [7] was the first introduced ray representation of solids. For a single direction and a uniform grid of rays parallel to that direction, the dexel structure stores the intervals of the rays lying inside the solid; these intervals are called *dexels* (depth elements). The G-buffer [26] extended the