Computers & Graphics 37 (2013) 669–675

Contents lists available at SciVerse ScienceDirect

Computers & Graphics

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

SMI 2013

Orthogonal slicing for additive manufacturing

Kristian Hildebrand [*], Bernd Bickel, Marc Alexa

TU Berlin, Einsteinufer 17, 10587 Berlin, Germany

A R T I C L E I N F O

Article history:
Received 29 March 2013
Received in revised form
24 May 2013
Accepted 27 May 2013
Available online 12 June 2013

Keywords:
Digital manufacturing
Shape analysis

A B S T R A C T

Most additive manufacturing technologies work by layering, i.e. slicing the shape and then generating each slice independently. This introduces an anisotropy into the process, often as different accuracies in the tangential and normal directions, but also in terms of other parameters such as build speed or tensile strength and strain. We model this as an anisotropic cubic element. Our approach then finds a compromise between modeling each part of the shape individually in the best possible direction and using one direction for the whole shape part. In particular, we compute an orthogonal basis and consider only the three basis vectors as slice normals (i.e. fabrication directions). Then we optimize a decomposition of the shape along this basis so that each part can be consistently sliced along one of the basis vectors. In simulation, we show that this approach is superior to slicing the whole shape in one direction, only. It also has clear benefits if the shape is larger than the build volume of the available equipment.

1. Introduction

Additive manufacturing techniques usually add layer after layer for fabricating a shape. Depending on the underlying process this introduces *direction bias*. The most obvious example for such bias is a different accuracy along the normal direction to a layer and the tangent directions. There are other factors that make the distinction of the directions worthwhile: different tensile strength or strain [1] (i.e. one can increase the stability of the model by choosing the right orientation in each part), different build time [2] (one can save production time by orienting the parts differently), different amounts of support material (i.e. one can save cost/waste by orientating different parts differently), or simply different dimensions of the build volume.

As a running example for our work we focus on the issue of accuracy. While our approach can be generalized to all layered manufacturing methods from 2D slabs laser cutting to high resolution 3D prints, we wish to stress that the improvements one can get from slicing one object into different directions may depend on its scale, the size of the object, and the desired application. The benefits of our method show in particular

- with increasing thickness of layers for laser cutting cardboard or plywood and low resolution 3D prints (i.e. high anisotropy of accuracy) or
- for large objects that cannot be fabricated as a whole because they do not fit the fabrication space.

Additionally multi-material objects that cannot be printed in one run because of the printer limitations or puzzles that are made to be of pieces for manually assembly are interesting applications for our method.

Our goal is to decompose the shape into few pieces so that each piece can be consistently sliced with small geometric error—and that by assembling the pieces one gets a replica with overall small error (see Fig. 1). The corresponding optimization problem needs to avoid both extremes: we assume that using one direction is not flexible enough, creates large error, or would simply be impossible; while decomposition into many pieces can clearly make the error small, but the assembly becomes tedious or virtually impossible.

One observation is that the previously mentioned directional bias introduces aliasing along the normal direction of the slice. This effect becomes more prominent with increasing layer thickness (e.g. >0.5 mm). When partitioning the object into several pieces using non-orthogonal aligned cuts and fabricating the parts along their optimal direction the assembly will not fit perfectly (resulting in a 'jaggy connection') and connections cannot be glued together properly. While orthogonality could be achieved locally for some cuts and assembly directions we suggest to solve this problem globally.

Our first modeling decision for this work is, consequently, to *restrict the slicing directions* as well as the normals of the cutting planes to an *orthogonal basis* $B=[\mathbf{b}_0\mathbf{b}_1\mathbf{b}_2], B^TB=I$. This approach allows selecting for each part independently an optimal slicing direction b_i while guaranteeing planar connection areas without sampling artifacts between parts (see Section 4).

---

[*] Corresponding author. Tel.: +49 3031428831.
E-mail address: kristian.hildebrand@tu-berlin.de (K. Hildebrand).

0097-8493/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cag.2013.05.011

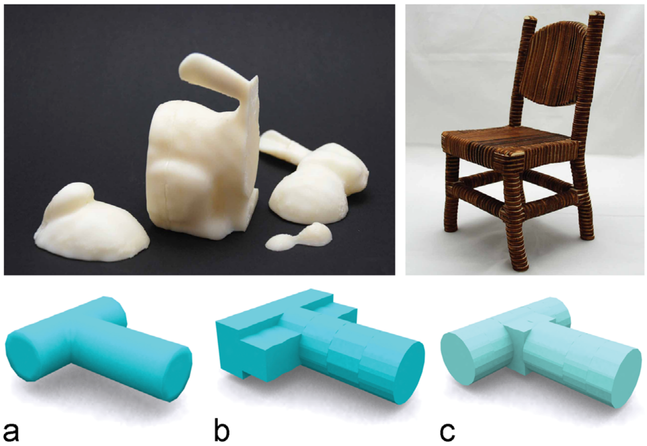

Fig. 1. Top: we present a framework that improves additive manufacturing methods across different scales: 3D printing resolutions (left) medium to large scales that might exceed the machine manufacturing volume (right). Bottom: (a) T-shaped object. (b) sliced in one direction (c) decomposed into two partitions after optimization and sliced in two directions.

We model the anisotropy in accuracy (or other properties in the process) as small cubic cells with dimension $d \times (d/N) \times (d/N)$, i.e. the thickness of a slice is d, while the accuracy in the tangent directions is N times better than the thickness of a slice. With this basic element, the most natural choice for a smallest element with consistent slicing direction is voxel cell of size d^3. Our idea is to pre-process the shape by decomposing it into voxels, and then find for each voxel its optimal slicing direction and a corresponding contour (see Section 5). We note that only voxels containing parts of the shape's boundary vary in their error depending on the direction.

With this information, we optimize a partition of the voxel set along the voxel faces. The goal is to generate large sets of voxels that are processed along the same direction.

2. Related work

Computer graphics and related fields in engineering have significantly contributed to computational approaches for computer-aided design that are essential tools in today's digital production pipeline. We will focus on a small subset of this work.

2.1. *Manufacturing and fabrication-oriented design*

Additive manufacturing methods are well evaluated and analyzed and show in various research approaches that optimization of the layered manufacturing process is essential. A number of methods address the task of finding an optimal orientation of a single part [3], considering surface finish, evaluate the surface roughness and part deposition time [4,5]. Danjou and colleagues [2] suggest an optimization procedure based on a genetic algorithm to improve the printing orientation. Masood et al. [6] show methodologies for computing the correct orientations based on the minimum volumetric error of basic primitives. Most closely related to our orientation optimization method, Reisner et al. [7] propose a method of finding an orthogonal frame. However, none of these approaches considers segmenting the model into subparts with different orientations.

In a broader context, Luo et al. [8] propose a segmentation algorithm to subdivide a mesh into pieces for the purpose of fitting a large model in a smaller 3D printing volume. This specifically focuses on finding structurally sound and aesthetic pleasing cutlines. In contrast, the goal of our work is to propose a framework to optimize the manufacturing process in accuracy.

By design, our method produces parts that can be simply glued together. There are a variety of approaches that generate specialized connectors used for furniture fabrication [9] or for connecting 3D printed parts [8]. We could easily incorporate this into our framework.

2.2. *Slicing and abstraction*

Planar elements play an essential role in shape analysis, approximation and abstraction. Sellamani et al. [10] gather prominent cross-sections that are also used for mesh segmentation. Recently, [11–13] proposed an approach for generating shape abstractions out of a minimal set of planar sections. Décoret et al. [14] use billboard clouds as an efficient shape representation. While these approaches rely on mostly unstructured sets of planar elements our proposed framework uses a regular set of stacked layers approximating the shape. Autodesk 123D [15] is able to create custom laser-cut sheets from a 3D shape. In contrast, this approach does not sufficiently take into account the orthogonal fabrication resolution and it is limited to one global slicing direction.

Slicing free-form surfaces was studied in the area of Computer Aided Design for example in the context of finding optimal milling machine paths [16,17]. Improving the geometric accuracy of layered manufacturing is proposed by Kulkani et al. [18].

3. Overview

Fig. 2 illustrates our pipeline to generate partitions that are sliced along good directions.

1. Given an input shape we compute a set of orthogonal directions $B = [\mathbf{b}_0 \mathbf{b}_1 \mathbf{b}_2]$ that are likely suited for a decomposition of the shape into small parts, each of which can be sliced along one of the three directions with small error (see Section 4). By rotation of the model with B^T we can now consider the canonical directions, i.e. x, y, z.
2. The shape is then decomposed into voxels of size d^3, where d is the desired slice thickness (or, worst accuracy). Each voxel is then decomposed into N^3 sub-voxels, where N is the factor between the thickness of the slice and the accuracy in the tangent directions.
3. For each voxel, the errors for each of the three slicing directions are computed. We use the discrete volumetric difference between the input shape and each of three approximations for a certain direction, computed on the sub-voxel grid. This requires computing approximations that are constant in the direction normal to a slice, yet may vary in tangent direction with the sub-voxel resolution. We explain how to do this consistently for all voxels, yet using only information available in each voxel in Section 5.
4. Based on the per-voxel errors, we compute a decomposition of the voxel grid so that each part can be sliced consistently with small error, yet the total number of pieces remains small. We also consider other factors in this process, such as the maximum size of each part. This process is explained in Section 6.

The result is an orthogonal decomposition of the shape into few pieces, as well as a direction for slicing for each piece.

4. Selection of manufacturing direction

Computing a set of orthogonal directions is the first step in our optimization. We base this computation on a simple observation: a