



ELSEVIER

Contents lists available at ScienceDirect

Computer Aided Geometric Design

www.elsevier.com/locate/cagd


Printing 3D objects with interlocking parts


 Peng Song^{a,*}, Zhongqi Fu^b, Ligang Liu^b, Chi-Wing Fu^c
^a School of Computer Science and Technology, University of Science and Technology of China, China

^b School of Mathematical Sciences, University of Science and Technology of China, China

^c School of Computer Engineering, Nanyang Technological University, Singapore

ARTICLE INFO

Article history:

Available online 30 March 2015

Keywords:

 Interlocking
 Shape partitioning
 Voxelization
 3D printing

ABSTRACT

Recent advances in 3D printing technologies bring wide range of applications from fast prototyping to product manufacturing. However, one intrinsic limitation of 3D printing is that we cannot fabricate a single object that is larger than the working volume of a 3D printer. To address this issue, we may partition the given object into 3D parts of manageable sizes for printing, and then assemble the object from the printed 3D parts. Rather than using connectors, glue, or skew, we propose to connect the printed 3D parts by 3D interlocking such that the assembled object can be not only repeatedly disassembled and reassembled, but also strongly connected by the parts' own geometry. To achieve these, we develop a voxelization-based approach to partition a given 3D model into 3D interlocking parts. To guarantee the generated 3D parts to be structurally sound and well-connected by 3D interlocking, we deform the local geometry of the 3D model to avoid voxel fragmentation, employ internal voxels to create initial interlocking parts, and analyze the local shape within voxels to guide the final parts construction. We demonstrate the effectiveness of our approach on 3D models with a variety of shapes, and realize some of them by 3D printing.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

3D printing is an additive manufacturing technology, facilitating convenient and rapid fabrication of physical objects of almost any shape. It has a wide range of practical applications, from fast product prototyping, product development, 3D visualization, to distributed manufacturing of larger-sized objects such as machine parts.

However, 3D printing has an intrinsic limitation: a 3D printer cannot directly print an object whose size is greater than the printer's working volume. This practical limitation has been pointed out recently by Luo et al. (2012), who proposed a solution to partition a given 3D object into parts for 3D printing and then assemble the printed parts together to reconstruct the object. Besides addressing the intrinsic limitation, this approach has several other advantages. First, it facilitates cost-effective maintenance since we only need to print a replacement part for a corresponding broken part rather than reprinting the entire object. Second, this approach is good for storage and transport since we can disassemble a large 3D object to save space and avoid breaking it during the transportation. Lastly, similar to Lego bricks, we could reassemble an existing 3D object, change some of its parts, and even reconfigure it for alternative designs and appearance.

* Corresponding author.

E-mail address: songpeng@ustc.edu.cn (P. Song).

To connect 3D printed parts, there are some common approaches, e.g., by male and female connectors (Luo et al., 2012; Lo et al., 2009) and by glue (Shapeways, 2014). For the case of male and female connectors, though they are generally practical, they may not provide sufficient structural strength to sustain the parts connections. Moreover, for these tiny printed features, they may be broken easily during the object assembly or the transportation. For the case of glue, though it could tightly join the printed 3D parts together, it is a permanent connection, discouraging object reassembly, cost-effective maintenance, and reconfiguration. In sharp contrast, we take a 3D interlocking approach (Xin et al., 2011; Song et al., 2012) to construct and connect printed 3D parts to form an object assembly. By this, we can overcome the above mentioned issues.

Connecting parts by 3D interlocking has several advantages: i) first, parts connections are achieved by the parts' own geometry without requiring the creation of extra tiny connectors such as male and female connectors; ii) the assembled object can be repeatedly disassembled and reassembled, facilitating cost-effective maintenance, storage, and transportation; iii) 3D interlocking is known to be strong, evidenced by its usage in long-standing architectural wooden structures; hence, this connection method allows us to achieve stronger 3D parts connections, which are in fact supported by inter-parts blockage with their geometry; and iv) lastly, it enables us to produce 3D parts with clean and smooth surface without hole drilling and protrusion.

In this paper, our goal is to partition a given 3D object into interlocking parts for 3D printing and object assembly. This is a highly challenging and unexplored problem, since we have to enforce not only the complicated 3D interlocking requirement (Xin et al., 2011; Song et al., 2012), but also the geometric and dimensional requirements on the printed 3D parts, as well as maintaining the 3D object appearance after its assembly. The requirements to be considered when we develop the computational method are listed as below:

- **3D interlocking:** the printed 3D parts should interlock one another, yet can be assembled and disassembled;
- **Printable parts:** while being not too small, each printed 3D part should fit into the working volume of the target 3D printer;
- **Structural soundness:** we should avoid thin and/or weak features on the 3D parts since such features could be easily broken during the object printing, assembly, or transportation;
- **Strong connection:** we should achieve strong 3D parts connections by ensuring strong blockage among the 3D parts in the 3D interlocking assembly; and
- **Aesthetics:** lastly, we should avoid having cutting seams that pass through salient regions on the assembled object surface since this affects the object appearance.

Note that state-of-the-art methods (Xin et al., 2011; Song et al., 2012) for creating 3D interlocking structures are insufficient to handle the problem since they mainly focus on the 3D interlocking requirement without considering the 3D printing issues such as structural soundness and object appearance.

To achieve the above requirements, we develop a novel voxelization-based approach to construct interlocking 3D parts from a given 3D model. Our technical contributions are three-fold. First, we develop a new framework that can create interlocking 3D parts from 3D models of general shape, where novel ideas include voxelizing the 3D model with local shape analysis, employing internal voxels to create initial interlocking parts, and attaching boundary voxels to initial parts while retaining the 3D interlocking. Second, we propose to deform the input model surface subject to the voxelization, which helps to avoid fragmented and disconnected shape features on the generated 3D parts. Lastly, we propose the shape and saliency connection graphs encoding local shape information to guide the parts geometric construction for achieving the parts structural soundness and aesthetics requirements. We have demonstrated the effectiveness of our approach on 3D models with a variety of shapes, and fabricated some of them by 3D printing to validate the parts connection capability, as well as their structural soundness.

2. Related work

3D fabrication. A number of research works in computer graphics have studied different aspects of 3D fabrication, e.g., deformation behavior (Skouras et al., 2013), mechanical characters (Coros et al., 2013; Ceylan et al., 2013), articulated models (Bächer et al., 2012; Cali et al., 2012), spinnable fabrication (Bächer et al., 2014), 3D shape balancing (Prévost et al., 2013), and printing material reduction (Wang et al., 2013; Lu et al., 2014).

Specifically, some of them focus on structural issues in 3D printing. Telea and Jalba (2011) devised several metrics to assess the printability of 3D shapes. Stava et al. (2012) proposed an automatic method to detect and correct structural issues in 3D models before printing them. Zhou et al. (2013) determined the worst-case loads in printing a 3D object. Umetani and Schmidt (2013) detected structural weakness of 3D objects for 3D printing optimization.

Fabrication by 3D parts. Partitioning an object into parts for 3D fabrication can be used for creating 3D furniture models with assemblable parts (Lau et al., 2011), for approximating objects with planar boundary pieces (Chen et al., 2013) or planar slices (McCrae et al., 2011; Hildebrand et al., 2012; Schwartzburg and Pauly, 2013), and for building architectural structures (Pottmann et al., 2007). Recently, Zhou et al. (2014) proposed a method to produce 3D printing objects that can be folded into a box.

In particular, some research works focus on partitioning a solid 3D shape into disjoint parts for 3D fabrication. Medellín et al. (2007) subdivided a 3D shape into parts based on a regular grid. Hao et al. (2011) decomposed a large complex model

into simpler 3D parts by using curvature-based partitioning. Luo et al. (2012) partitioned a large 3D object into smaller parts by using planer cuts, where male and female connectors are generated on parts for making the connections. Hildebrand et al. (2013) addressed the directional bias issue in 3D printing by segmenting a 3D model into a few parts and assigning an optimal printing orientation for each part. Vanek et al. (2014) improved the 3D printing efficiency by converting 3D objects into shells and breaking them into parts that can be glued together.

Our work also partitions a 3D shape into smaller parts for 3D fabrication. But in sharp contrast, we achieve 3D parts connections by 3D interlocking rather than by glue or by tiny geometry like male and female connectors. Our approach has several distinctive advantages over existing approaches as highlighted earlier in the introduction.

3D interlocking puzzles. Designing 3D interlocking puzzles is extremely difficult even for highly skilled puzzle designers (Coffin, 1990). So far, there are only a few computational methods that can construct 3D interlocking structures. Xin et al. (2011) created interlocking puzzles from 3D models by replicating and connecting a specific six-piece burr configuration, but since this method requires a centralized bulky structure to achieve 3D interlocking, it is highly restrictive, and could not deal with complex shapes and general topology. Later, Song et al. (2012) developed a recursive method to construct 3D interlocking structures by iteratively extracting puzzle pieces from a voxelized model. The produced puzzle pieces are polycubes with cubical appearance, and cutting seams are randomly arranged on the object surface, resulting in rather distracting appearance. If we directly apply constructive solid geometry (CSG) operations to refine their puzzle pieces with the original object surface, disconnected, fragmented, and weakly-connected components would be easily produced on the puzzle pieces, making (Song et al., 2012) inadequate for 3D printing, see Section 6 for detailed comparison.

Compared with the above works, our method enforces various fabrication requirements while achieving 3D interlocking, which is already nontrivial on its own. In particular, our method not only can automatically and flexibly create interlocking 3D parts from a given object model of general shape, but also can produce 3D parts that are structural sound and avoid obvious cutting seams on salient features of the object. By this, a given object is ready for being 3D printed with its parts.

3. Overview

In this section, we first discuss the challenges of partitioning a given 3D shape into interlocking 3D parts, and then give an overview of our approach.

Challenges of 3D shape partitioning. For a given 3D shape, there exist different ways of partitioning it into parts, for example, plane – (Luo et al., 2012), curvature – (Hao et al., 2011), and voxelization-based (Medellín et al., 2007; Zhou et al., 2014) approaches. This work employs the voxelization-based partitioning approach to create interlocking parts since i) cubical voxel space facilitates the design of interlocking structures like (Song et al., 2012); ii) shape analysis of parts can be simplified by computing the local shape within voxels and connecting local shapes between neighboring voxels to form component parts; and iii) parts dimension can be easily controlled by the bounding box in voxel space.

However, there are two major challenges to construct interlocking parts by voxelization. First, when voxelizing a 3D model, a voxel may be full, partial, or empty according to the local shape within the voxel, see Fig. 1(a). More importantly, a resulting partial voxel may have various kinds of local shapes, e.g., a tiny fragment (Fig. 1(b) (case 1)), disconnected fragments within a voxel (Fig. 1(b) (case 2)), or a thin structure (Fig. 1(b) (case 3)). Hence, a naive voxelization could easily produce a large amount of partial voxels with undesired (e.g., tiny or disconnected) local shapes. Second, when we later connect neighboring voxels to form 3D component parts in the assembly, some problematic parts can be generated if we arbitrarily connect neighboring voxels into parts, e.g., parts with a weak fragment (see P1 in Fig. 1(c)) and parts with disconnected fragments (see P2 in Fig. 1(c)). For the case of P2, we may not just solve it by separately attaching the two fragments to different parts since this may violate the 3D interlocking.

To address the above issues, we propose to analyze the local shape within each partial voxel, and apply the information to guide the voxelization process as well as the 3D parts construction. By this, we can achieve structurally-sound component parts (see P3 in Fig. 1(c)) in the 3D interlocking assembly.

Overview of our approach. We take a 3D watertight surface as input since only meshes with well-defined exterior boundary and closed volume can be 3D printed, see Fig. 2(a). Note that an input mesh that does not satisfy this requirement (e.g., with self-intersections) can be repaired by methods such as (Attene, 2014).

Given an input watertight surface, we first place a voxel grid in the 3D object space and slightly adjust the voxel size to reduce the number of undesired partial voxels. Then, we voxelize the input model. For partial voxels with tiny fragment(s), we avoid the fragment(s) by slightly deforming the model geometry locally. Next, we measure the connection strength between neighboring (partial) voxels and build a shape connection graph by analyzing the local shapes (Fig. 2(b)). We also identify salient appearance features on the model by (Lee et al., 2005), and represent the information by a saliency connection graph.

After the above preparation works, we construct an initial set of interlocking 3D parts from the internal voxels (Fig. 2(c)). Hence, the initial 3D parts can fulfill the structural soundness and strong connection requirements since all employed

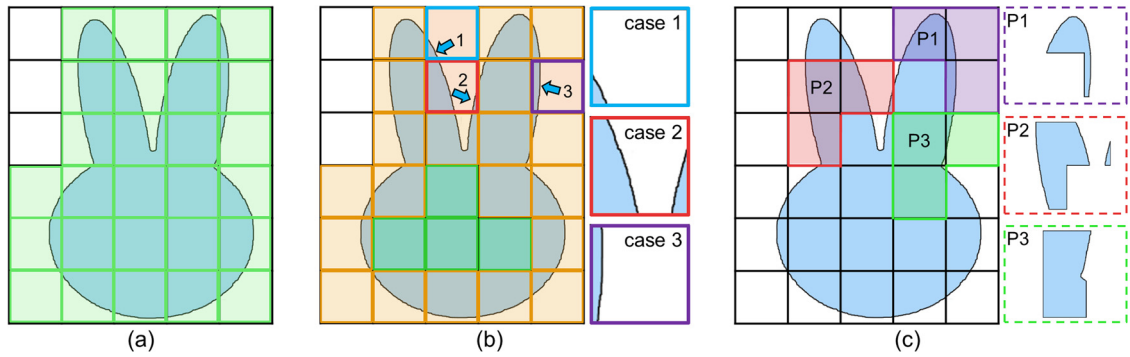


Fig. 1. Voxelization-based shape partitioning. (a) A 2D shape and its voxelization; (b) full (in green) and partial (in orange/blue/red/purple) voxels: the three partial voxels in blue, red, and purple show cases of tiny fragments (case 1), disconnected fragments (case 2), and thin structures (case 3), respectively; (c) naively connecting voxels into component parts could easily produce thin and disconnected 3D parts (in purple and red) rather than structurally-sound 3D parts (in green). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

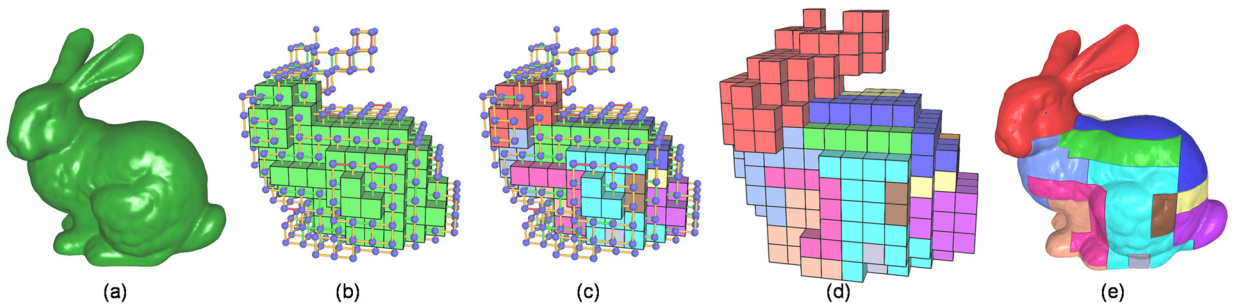


Fig. 2. Overview. (a) A 3D watertight model as input; (b) voxelization and shape analysis; (c) creating initial interlocking parts; (d) attaching boundary voxels to the parts; and (e) refining the parts' geometry by CSG intersection with the input mesh model.

voxels are large enough and thus provide sufficient blockage to enforce the 3D interlocking. After that, we attach the remaining partial voxels one by one to the constructed interlocking parts as guided by the shape connection graph (Fig. 2(d)) without breaking the fulfilled requirements. Later, we enforce the aesthetics requirement by reassigning boundary voxels among neighboring parts as guided by the saliency connection graph, aiming to avoid cutting seams on salient object features.

Lastly, we construct the geometry of the final 3D parts (Fig. 2(e)) by using CSG intersection between the voxelized parts (Fig. 2(d)) and the input mesh (Fig. 2(a)).

4. Voxelization and shape analysis

This section presents the various preparation works we have before constructing the geometry of interlocking 3D parts from an input 3D model. This includes voxelization, local shape deformation, and the generation of shape connection graph and saliency connection graph.

Voxelization method. Given the input watertight 3D model, we first voxelize it by (Nooruddin and Turk, 2003), which casts parallel rays through the model and uses parity count to classify voxels as interior or exterior. We improve this method by casting additional parallel rays to not only classify voxels but also estimate the local shape within each voxel. For a bounded volume of $W \times H \times D$ voxels, we first build a $(K \times W + 1) \times (K \times H + 1) \times (K \times D + 1)$ uniform 3D point grid within it, where we have $K + 1$ sample points along the edge of each voxel. Then we cast $(K \times W + 1) \times (K \times H + 1)$ rays through the model to classify each sample point as interior or exterior, where each ray passes through $(K \times D + 1)$ sample points. Note that we can do ray casting along any of the three major axis of the grid. Fig. 3 illustrates our voxelization method with a 2D example having $K = 6$.

Next, we calculate properties of the local shape within each voxel based on the classified sample points:

- *Local shape volume* is estimated by counting the number of interior sample points and then computing their coverage percentage within the voxel;
- *Type of the voxel* is classified as full, partial, or empty based on the value of local shape volume, i.e., 1, (0, 1), and 0, respectively;
- *Contacting face area* for a partial voxel is estimated by computing the percentage of interior sample points on the voxel's contacting face with its neighbor, see the yellowish-green sample points in Fig. 3(b);

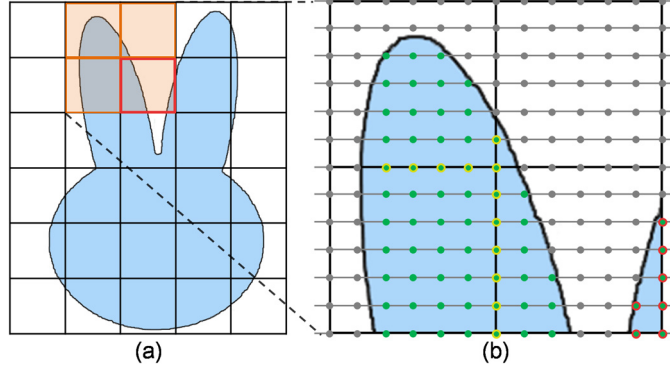


Fig. 3. (a) We take a 2D shape to illustrate our voxelization method; (b) classified interior and exterior sample points are in green and gray, respectively; the bottom-right voxel is detected as a disconnected partial voxel by breadth first search from an interior sample point (see red circles). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 4. Deform the model's local geometry (a & b) to remove a partial voxel with a tiny fragment, and (c & d) to resolve a partial voxel with disconnected local shape.

- *Local shape connectivity* in a partial voxel is checked by i) breadth first search from an interior sample point, and ii) checking whether it reaches all interior sample points in the voxel.

To facilitate shape partitioning, partial voxels with tiny or disconnected local shape should be avoided, see again Fig. 1. Hence, for voxel grid with a given resolution, we first adjust the voxel size slightly to reduce the number of undesired partial voxels. This is achieved by randomly sampling the voxel size within a small range and selecting the voxel size corresponding to the least number of undesired voxels. Note that we do not adjust the 3D model pose (orientation) to avoid undesired partial voxels since this could result in non-axis-aligned cuts on the model, affecting both object aesthetics and parts assembly.

Local deformation to avoid tiny fragments. For undesired partial voxels in the voxelization, we propose to resolve them by locally deforming the 3D model geometry. To constrain the local deformation within an undesired voxel and its neighbors, we remesh the 3D model (Alliez et al., 2008; Botsch et al., 2010) such that the mesh triangles are regular and much smaller than the voxel size. Note that we can detect tiny fragments by looking at partial voxels whose local shape volume is smaller than a threshold, which is set as 0.01 in our experiments.

For a partial voxel with a tiny fragment (Fig. 4(a)), we compute three variables for the fragment by examining (averaging) the mesh triangles related to the tiny fragment in the voxel: its centroid V_c , mean surface normal \vec{N} , and radius of a bounding sphere r . For a partial voxel with disconnected tiny fragments (Fig. 4(c)), the same three variables are computed for each tiny fragment. To avoid a tiny fragment, we could locally deform the 3D model geometry by iteratively moving the related mesh vertices by using

$$V_i^{k+1} = V_i^k - \alpha (2r - \text{dist}(V_i^k, V_c)) \vec{N}, \quad (1)$$

where V_i is a mesh vertex within a bounding sphere centered at V_c with radius $2r$, and α is a parameter to control the deformation, where we put α as 0.01 in our implementation. Note that we have to recompute the three variables (i.e., V_c , \vec{N} , and r) for the deformed tiny fragment in each iteration, and the iterative process terminates when the tiny fragment vanishes in the related voxel, see Fig. 4(b & d).

Internal and boundary voxels. Partial voxels mostly cannot provide sufficient blocking strength for enforcing interlocking. Hence, one key strategy to achieve parts interlocking is by *considering full voxels to produce interlocking parts and later attaching partial voxels to the parts*. However, for most 3D models, the number of full voxels could be rather small, thus restricting the construction of interlocking parts, especially when the voxel grid has low resolution.




Fig. 5. Voxelizing Duck. (a) Full voxels in the model; (b) internal voxels identified; (c) voxelized model; and (d) shape connection graph for the boundary voxels.




Fig. 6. (a) Internal (green) and boundary voxels; (b) shape connection graph for the boundary voxels, where none, weak, normal, and strong connections are indicated in gray, red, yellow, and green, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 7. (a) The saliency map on Duck; (b) the saliency connection graph with the same color coding as Fig. 6(b).

To include more voxels for achieving the interlocking, we also take some partial voxels into account based on the following constraints: i) large local shape volume; ii) no voxel face with tiny contacting area; and iii) has at least one full voxel as its neighbor. We call such partial voxels and the full voxels the *internal voxels*, which form the *internal volume*, and we call the remaining partial voxels outside the internal volume the *boundary voxels*, see Fig. 5.

Shape connection graph. For the boundary voxels, we build shape connection graph to describe their structural connection strength with neighboring voxels, see Fig. 5(d) and Fig. 6(b) for examples. Such connection can be identified as none, weak, normal, or strong depending on the normalized contacting face area we computed earlier (in our experiments, the values are set to be 0, (0, 0.05], (0.05, 0.25], [0.25, 1.0] respectively). The shape connection graph will later help to guide the attachment of boundary voxels to initial interlocking parts, aiming to avoid the creation of disconnected or fragile 3D parts.

Saliency connection graph. In addition, we build saliency connection graph for boundary voxels to later avoid putting cutting seams (between interlocking 3D parts) on salient object features. In detail, we first apply (Lee et al., 2005) to compute the saliency value at every mesh vertex (see Fig. 7(a)), and then estimate the saliency value of each boundary voxel by averaging the saliency values of all mesh vertices it contains. The saliency connection strength between two neighboring boundary voxels is defined as the product of the two voxels' saliency values, since we have a strong saliency connection only when the two related voxels have large saliency values, see Fig. 7(b) for an example.

5. Generating 3D interlocking parts

Given the various inputs we prepared in Section 4, there are four major steps in generating 3D interlocking parts:

Step 1: Generating initial 3D interlocking parts. First, taking the voxels in the internal volume as input, we employ (Song et al., 2012) to construct an initial set of interlocking 3D parts, see Fig. 2(c). Here we give a short description of the randomized algorithm in (Song et al., 2012), which takes a general voxelized shape as input and iteratively extracts puzzle pieces from it to create interlocking puzzles. A formal model is proposed in (Song et al., 2012) to ensure global interlocking of all extracted puzzle pieces by enforcing local interlocking requirement on every three consecutive intermediate puzzle pieces. To achieve this local interlocking requirement, the geometry of each puzzle piece is constructed based on a family of strategies that manipulate voxel blocking mechanics.

Here we revise (Song et al., 2012) in two aspects:

- *General interlocking.* We generalize the formal model in (Song et al., 2012) to support the creation of both general interlocking and recursive interlocking structures such that we can generate interlocking parts from a given 3D model more flexibly. That is, the extracted 3D parts are general interlocking (recursive interlocking) if we enforce local interlocking requirement on every $K > 3$ ($K = 3$) consecutive intermediate parts when constructing them.
- *Parts' size.* To meet the printable parts requirement, the size of each 3D part should not exceed the 3D printer's working volume. Hence, when constructing a 3D part, we keep updating its bounding box when it grows according to the procedure in (Song et al., 2012), and we attach a new voxel to it only if its bounding box is still contained within the 3D printer's working volume.

By the above considerations, we can enforce the 3D interlocking and printable parts requirements for the initial 3D parts. Moreover, since the internal voxels involved here do not have tiny/thin fragments, and they have large local shape volume to support 3D parts blocking, they can help to fulfill the structural soundness and strong connection requirements for the initial 3D parts.

Step 2: Attaching boundary voxels to initial 3D parts. To assemble back the given 3D object, all boundary voxels have to be attached to the initial 3D parts while maintaining the various fabrication requirements. Note that attaching voxels to the initial 3D parts could strengthen the parts blocking, so the strong connection requirement can always be maintained.

We develop an iterative method to attach boundary voxels one by one to the initial 3D parts, see Algorithm 1 for the details. Note that for a given set of n interlocking 3D parts, we assume the disassembly order to be P_1, P_2, \dots , until P_n , with P_1 as the key part that locks the entire assembly structure. For a given boundary voxel, we try to find one of its neighboring parts to attach the voxel, without violating the fabrication requirements. Here we employ the shape connection graph to ensure a boundary voxel is attached to a 3D part only when they share strong structural connection (i.e., large contacting area).

In Algorithm 1, we need to perform the voxel attachment (i.e., the 1st *for loop*) for several iterations (i.e., the *while loop*) such that the boundary voxels that are failed to be attached in current iteration could be attached in the next iteration. This is because a boundary voxel can have more choices of parts to be attached to after its neighboring boundary voxels have been assigned to certain parts. Due to the many constraints on the voxel attachment, our algorithm does not guarantee all boundary voxels can be attached, especially for 3D models with complex shapes such as the Buddha in Fig. 9. In such case, we could regenerate initial 3D parts and then repeat Algorithm 1.

Step 3: Cutting seam refinement. We further enforce the aesthetics requirement by reassigning voxels among neighboring parts to avoid putting cutting seams across salient object features. We first identify pairs of neighboring boundary voxels that have strong saliency connection (using the saliency connection graph) but belong to different 3D parts. Then, we try to put each boundary voxel pair into the same 3D part by searching for possible ways to swap voxels among different 3D parts within the neighborhood. During such voxel reassignment, we have to maintain the interlocking requirement for the modified 3D parts, as well as avoid structural issue. Fig. 8(b & c) shows an example of refining cutting seams by reassigning voxels as guided by the mesh saliency (see Fig. 8(a & d)).

Step 4: Parts surface refinement. To account for the machine fabrication tolerance, there should be sufficient empty space in-between 3D parts. Hence, we create the empty spacing by thinning the voxels on each part according to the voxels' local connectivity with their neighbors. The final parts' geometry can then be obtained by CSG intersection between each voxelized part and the original 3D model, see Fig. 8. Comparing Fig. 8(e & f), we can see that our voxel-reassigning strategy can help avoid undesirable cutting seams on the assembled object surface.

6. Results

Our method enables us to create 3D interlocking parts from object models of various shapes, see Fig. 9. All the generated 3D parts are well-connected by interlocking upon the parts assembly, see Fig. 10 for three more examples.

Algorithm 1: Attaching boundary voxels to initial 3D parts.

Input: Initial interlocking 3D parts;
All the boundary voxels.

Output: Flag to indicate if the attachment succeeds;
Modified parts with boundary voxels attached.

```

Put all boundary voxels into a set  $S$  ;
while  $S$  is non-empty do
  Initialize number of attached boundary voxel  $m = 0$  ;
  for each boundary voxel  $V_i$  in  $S$  do
    Find all the neighboring parts of  $V_i$  ;
    for each neighboring part  $P_k$  do
      if there is no or weak shape connection between  $V_i$  and  $P_k$  (using the shape
      connection graph) then
        | continue ;
      end
      if attaching  $V_i$  to  $P_k$  violates the dimension constraint of  $P_k$  then
        | continue ;
      end
      if attaching  $V_i$  to  $P_k$  blocks the movement of any piece from  $P_1$  to  $P_k$  then
        | continue ;
      end
      Push  $P_k$  into a candidate part vector  $S_P$  ;
    end
    if  $S_P$  is non-empty then
      | Select a part  $P_s$  in  $S_P$  which has the largest contacting area with  $V_i$  ;
      | Attach  $V_i$  to  $P_s$  ;
      | Remove  $V_i$  from  $S$  ;
      |  $m++$  ;
    end
  end
  if  $m$  is equal to 0 then
    | Return false ;
  end
end
Return true ;

```

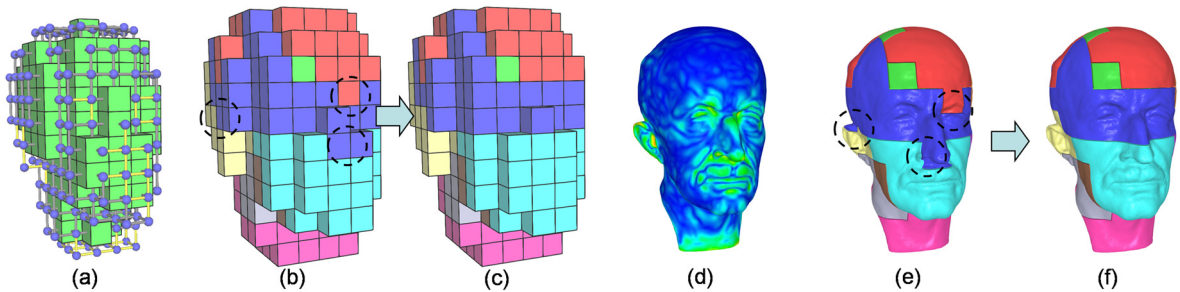


Fig. 8. (a) Saliency connection graph for MAXPLANCK; (b & c) voxelized 3D parts before and after cutting seam refinement; (d) saliency map on MAXPLANCK; (e & f) CSG-refined 3D parts without and with cutting seam refinement. Note that the reassigned voxels and undesirable cutting seams are marked by black circles in (b) and (e).



Fig. 9. Various 3D interlocking models created by our approach: 8-parts BIMBA, 10-parts MICKEY, 16-parts RABBIT, 10-parts VASE, 16-parts PUMPKIN, and 20-parts BUDDHA (from left to right).




Fig. 10. Assembly of a 10-parts BALL (top), a 12-parts TURTLE (middle), and a 12-parts DUCK (bottom).

Table 1

Timing statistics and printing volume reduction.

	Timing statistics				Saving printing volume		
	# Model vertices	Volume resolution	# Parts	Timing (s)	Max model length	Max part length	Length ratio
Duck	3290	$9 \times 10 \times 11$	6	25	2.00	1.33	66.5%
Bimba	15 653	$10 \times 11 \times 7$	8	42	2.00	1.27	63.5%
Maxplanck	49 132	$7 \times 11 \times 8$	10	62	2.00	1.23	61.5%
Bunny	15 002	$12 \times 12 \times 9$	12	43	2.00	1.51	75.5%
Rabbit	10 139	$7 \times 16 \times 8$	16	32	2.00	0.89	44.5%
Buddha	44 973	$11 \times 16 \times 11$	20	150	2.00	1.16	58.0%

Implementation and performance. We implemented our method in C++ and ran our experiments on a desktop computer with a 3.4 GHz CPU and 8GB memory. For a given 3D model, users can specify the number of parts to be generated. Although our method allows the creation of a large number of 3D parts from a given model, we suggest partitioning the model into around 10 parts since a larger number of parts require more time for 3D printing and physical assembly. In some situations, we disable certain fabrication requirements to speed up the procedure of generating parts: i) we can disable the printable parts requirement for small objects, e.g., the ring in Fig. 13; and ii) we can disable the aesthetics requirement for object without salient features, e.g., the sphere in Fig. 10.

Table 1 presents the performance of our method. From the timing statistics shown on the 5th column from the left, we can see that our method can generate interlocking parts fairly efficiently. Note that such time performance depends on several factors: the number of 3D parts, the complexity of the model's shape, the resolution of the input mesh, and the resolution of the voxelization. Columns 6 to 8 in the table estimate how much printing volume could be reduced after partitioning an input model into six to twenty parts. Such performance depends not only on the number of 3D parts but also on the shape complexity, since complex shapes impose stronger constraints on the parts structure and aesthetics, which may conflict with the parts dimension constraint. For example, considering the 20-parts BUDDHA (see Fig. 9 and Table 1), the printing volume reduction is not that obvious since there are many salient details on its surface.

Table 2 shows the performance of our voxelization method in resolving partial voxels with tiny or disconnected fragments, as well as in considering partial voxels for creating the interlocking. Columns 3 and 4 from the left show that in the initial voxelization, there are a few voxels with disconnected fragments (i.e., disconnected voxels) and several voxels with tiny fragments (i.e., tiny voxels). By applying local deformation on the 3D model, disconnected voxels can be avoided in the voxelization (i.e., either be removed or become a connected voxel) while tiny voxels can also be successfully removed. In addition, columns 6 and 7 in the table show that our voxelization method can effectively include a large amount of partial voxels to construct a larger internal volume for creating initial 3D interlocking parts.

Comparison with Song et al. (2012). To apply (Song et al., 2012) for 3D printing, we need to refine the interlocking polycubes generated from (Song et al., 2012) by performing CSG intersection with the model surface. However, this straightforward solution has several problems. First, Song et al. (2012) do not consider the structural constraint when connecting voxels to form 3D parts (see again Fig. 1), so their generated 3D parts could have tiny and/or thin shape features after the CSG

Table 2

Voxel statistics of our voxelization method.

	Volume resolution	Before deformation			After deformation			
		# Disconnected voxels	# Tiny voxels	# Total voxels	# Full voxels	# Internal voxels	# Boundary voxels	# Total voxels
Duck	$9 \times 10 \times 11$	2	16	547	146	270	260	530
Bimba	$10 \times 11 \times 7$	3	15	425	75	206	202	408
Maxplanck	$7 \times 11 \times 8$	1	18	364	80	170	176	346
Bunny	$12 \times 12 \times 9$	4	32	536	116	224	276	500
Rabbit	$7 \times 16 \times 8$	0	17	533	123	260	256	516
Buddha	$11 \times 16 \times 11$	24	25	1096	381	665	402	1067




Fig. 11. Duck (12 parts) created by (Song et al., 2012) (top) and our method (bottom). (a & d) Voxelized 3D parts assembly; (b & e) CSG-refined parts assembly; and (c & f) three example 3D parts.

surface refinement. Second, some of these 3D parts may even be disconnected, see the red piece in Fig. 11(c). Third, since Song et al. (2012) do not consider the aesthetics constraint, it could generate a lot of distracting cutting seams across salient features on the object assembly, see Fig. 11(b). Through the local shape analysis procedure, our method can avoid putting cutting seams on salient object features, see Fig. 11(e), as well as avoiding the generation of tiny, thin, and disconnected shape features on the 3D parts, see Fig. 11(f).

Printed examples. We use 3D printers of three different types of printing technology to fabricate our interlocking 3D parts: Fused Deposition Modeling (FDM) with PLA plastic material (MakerBot Replicator 2 (Makerbot, 2014)), Selective Laser Sintering (SLS) with nylon material (Sinterstation HiQ SLS (3D Systems, 2014)), and Stereolithography (SLA) with Visijet Proplast (Projet 3500 HDMax (3D Systems, 2014)).

To validate the effectiveness of our method for printing a large 3D object with parts, we fabricate a 6-parts Duck model ($30 \text{ cm} \times 27 \text{ cm} \times 25 \text{ cm}$) using a DIY FDM 3D printer of working volume $25 \text{ cm} \times 25 \text{ cm} \times 45 \text{ cm}$. Since the object is too large for the printer's working volume, we partition it into 6 interlocking 3D parts with smaller sizes (Fig. 12(a)), among which the maximum length is only 20 cm. Hence, we can fabricate each part by our 3D printer. To validate the connection capability by 3D interlocking, we assemble the printed parts of various models (Fig. 12), and shake the parts assemblies. The parts assemblies are found to be steady and stable upon shaking, see the supplementary video.

Besides the Duck, we fabricate a 12-parts BUNNY (Fig. 12(b)) and a 10-parts VASE with four different colors (Fig. 12(c)) using the MakerBot FDM printer, from which we can see that the generated 3D parts are strong enough for printing and assembly. Note that this FDM printer can only print with a maximum resolution of 0.2 mm. To fabricate models with details, we print the 10-parts MAXPLANCK model (Fig. 12(d)) by using a 3D Systems SLS printer with 0.06 mm printing resolution, so that we can achieve smoother appearance on the assembled model.

Interlocking puzzle ring. Lastly, we apply our method to create a 22-parts interlocking ring puzzle shown in Fig. 13. It has a diameter of 3.2 cm, and we fabricate its parts by a 3D Systems SLA printer with 0.025 mm printing resolution. The assembled ring can be worn on a finger. We specially design its key piece: to take it out from the puzzle, we need to first lift it up for a short distance and then move it toward the ring center (see the supplementary video). Hence, when we wear this puzzle ring on our finger, our finger can block the movement of the key piece, thus preventing the ring from




Fig. 12. 3D-printed interlocking parts (left) and the corresponding assembled object (right). (a) 6-parts DUCK by FDM; (b) 12-parts BUNNY by FDM; (c) 10-parts VASE by FDM; and (d) 10-parts MAXPLANCK by SLS.




Fig. 13. The 22-parts interlocking ring created by our method. (a) Virtual puzzle ring model; (b) the fabricated (assembled) ring model; (c) the twenty two 3D parts, individually printed by SLA; and (d) the ring being worn on a finger.

disassembly. Note that such interlocking puzzle could be easier to play with than those from Song et al. (2012) since object features on the 3D parts could give hints to the puzzle assembly.

7. Conclusion

This paper presents a novel voxelization-based method to partition a given 3D object model of general shape into interlocking 3D parts such that the 3D object can be printed with smaller parts that are connected by 3D interlocking. We first voxelize the input 3D model and analyze the local shape within each voxel, where a local deformation strategy is developed to avoid voxel fragmentation in the voxelization. Second, we differentiate internal and boundary voxels according to the local shape contained and employ internal voxels to create initial 3D interlocking parts. Lastly, we apply the local shape information encoded by shape and saliency connection graphs to guide the construction of final parts' geometry. By these, we can generate 3D parts that are structurally sound and strongly connected by 3D interlocking, without obvious cutting seams across salient object features. We validate the performance of our method by fabricating a variety of models using different types of 3D printers, and also demonstrate its feasibility by designing and making the 22-parts interlocking ring.

Limitations. First, our method cannot generate interlocking 3D parts from a hollowed object with thin boundary since we need a solid volume to create parts blocking. Second, our method cannot control the positions of cutting seams on the input model and some of the cuts will be visible after the parts assembly. Third, the original shape of the input model could be changed slightly in the local deformation, so our current results may not be suitable for applications that require high model precision. Lastly, our shape analysis method cannot detect weak connections locally in a voxel; this issue could be circumvented by (Zhou et al., 2013), or alleviated by increasing the voxel grid resolution.

Acknowledgements

The authors would like to thank Yael Friedman, a jewelry designer, for her various helpful comments on designing the 22-parts interlocking ring. The project is partially supported by the National Natural Science Foundation of China (61403357,

61222206), the Fundamental Research Funds for the Central Universities (WK0110000044), One Hundred Talent Project of the Chinese Academy of Sciences, and MOE AcRF Tier2 funding (MOE2011-T2-2-041 (ARC 5/12)), Singapore.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <http://dx.doi.org/10.1016/j.cagd.2015.03.020>.

References

- 3D Systems, 2014. <http://www.3dsystems.com/>.
- Alliez, P., Ucelli, G., Gotsman, C., Attene, M., 2008. Recent advances in remeshing of surfaces. In: *Shape Analysis and Structuring, Mathematics and Visualization*, pp. 53–82.
- Attene, M., 2014. Direct repair of self-intersecting meshes. *Graph. Models* 76 (6), 658–668.
- Bächer, M., Bickel, B., James, D.L., Pfister, H., 2012. Fabricating articulated characters from skinned meshes. *ACM Trans. Graph. (SIGGRAPH)* 31 (4), article 47.
- Bächer, M., Whiting, E., Bickel, B., Sorkine-Hornung, O., 2014. Spin-It: optimizing moment of inertia for spinnable objects. *ACM Trans. Graph. (SIGGRAPH)* 33 (4), article 96.
- Botsch, M., Kobbelt, L., Pauly, M., Alliez, P., Lévy, B., 2010. *Polygon Mesh Processing*. AK Peters/CRC Press.
- Cali, J., Calian, D.A., Amati, C., Kleinberger, R., Steed, A., Kautz, J., Weyrich, T., 2012. 3D-printing of non-assembly, articulated models. *ACM Trans. Graph. (SIGGRAPH Asia)* 31 (6), article 130.
- Ceylan, D., Li, W., Mitra, N.J., Agrawala, M., Pauly, M., 2013. Designing and fabricating mechanical automata from mocap sequences. *ACM Trans. Graph. (SIGGRAPH Asia)* 32 (6), article 186.
- Chen, D., Sitthi-amorn, P., Lan, J.T., Matusik, W., 2013. Computing and fabricating multiplanar models. *Comput. Graph. Forum (EuroGraphics)* 32 (2), 305–315.
- Coffin, S.T., 1990. *The Puzzling World of Polyhedral Dissections*. Oxford University Press.
- Coros, S., Thomaszewski, B., Noris, G., Sueda, S., Forberg, M., Sumner, R.W., Matusik, W., Bickel, B., 2013. Computational design of mechanical characters. *ACM Trans. Graph. (SIGGRAPH)* 32 (4), article 83.
- Hao, J., Fang, L., Williams, R.E., 2011. An efficient curvature-based partitioning of large-scale stl models. *Rapid Prototyping J.* 17 (2), 116–127.
- Hildebrand, K., Bickel, B., Alexa, M., 2012. crdbrd: shape fabrication by sliding planar slices. *Comput. Graph. Forum (EuroGraphics)* 31 (2), 583–592.
- Hildebrand, K., Bickel, B., Alexa, M., 2013. Orthogonal slicing for additive manufacturing. *Comput. Graph. (SMI)* 37 (6), 669–675.
- Lau, M., Ohgawara, A., Mitani, J., Igarashi, T., 2011. Converting 3D furniture models to fabricatable parts and connectors. *ACM Trans. Graph. (SIGGRAPH)* 30 (4), article 85.
- Lee, C.H., Varshney, A., Jacobs, D.W., 2005. Mesh saliency. *ACM Trans. Graph. (SIGGRAPH)* 24 (3), 659–666.
- Lo, K.-Y., Fu, C.-W., Li, H., 2009. 3D polyomino puzzle. *ACM Trans. Graph. (SIGGRAPH Asia)* 28 (5), article 157.
- Lu, L., Sharf, A., Zhao, H., Wei, Y., Fan, Q., Chen, X., Savoye, Y., Tu, C., Cohen-Or, D., Chen, B., 2014. Build-to-last: strength to weight 3D printed objects. *ACM Trans. Graph. (SIGGRAPH)* 33 (4), article 97.
- Luo, L., Baran, I., Rusinkiewicz, S., Matusik, W., 2012. Chopper: partitioning models into 3D-printable parts. *ACM Trans. Graph. (SIGGRAPH Asia)* 31 (6), article 129.
- Makerbot, 2014. <http://www.makerbot.com/>.
- McCrae, J., Singh, K., Mitra, N.J., 2011. Slices: a shape-proxy based on planar sections. *ACM Trans. Graph. (SIGGRAPH Asia)* 30 (6), article 168.
- Medellín, H., Lim, T., Corney, J., Ritchie, J.M., Davies, J.B.C., 2007. Automatic subdivision and refinement of large components for rapid prototyping production. *J. Comput. Inf. Sci. Eng.* 7 (3), 249–258.
- Nooruddin, F.S., Turk, G., 2003. Simplification and repair of polygonal models using volumetric techniques. *IEEE Trans. Vis. Comput. Graph.* 9 (2), 191–205.
- Pottmann, H., Liu, Y., Wallner, J., Bobenko, A., Wang, W., 2007. Geometry of multi-layer freeform structures for architecture. *ACM Trans. Graph. (SIGGRAPH)* 26 (3), article 65.
- Prévost, R., Whiting, E., Lefebvre, S., Sorkine-Hornung, O., 2013. Make it stand: balancing shapes for 3D fabrication. *ACM Trans. Graph. (SIGGRAPH)* 32 (4), article 81.
- Schwartzburg, Y., Pauly, M., 2013. Fabrication-aware design with intersecting planar pieces. *Comput. Graph. Forum (EuroGraphics)* 32 (2), 317–326.
- Shapeways, 2014. *Gluing 3D printed parts*. http://www.shapeways.com/tutorials/gluing_3d_printed_parts_tutorial.
- Skouras, M., Thomaszewski, B., Coros, S., Bickel, B., Gross, M., 2013. Computational design of actuated deformable characters. *ACM Trans. Graph. (SIGGRAPH)* 32 (4), article 82.
- Song, P., Fu, C.-W., Cohen-Or, D., 2012. Recursive interlocking puzzles. *ACM Trans. Graph. (SIGGRAPH Asia)* 31 (6), article 128.
- Stava, O., Vanek, J., Benes, B., Carr, N., Měch, R., 2012. Stress relief: improving structural strength of 3D printable objects. *ACM Trans. Graph. (SIGGRAPH)* 31 (4), article 48.
- Telea, A., Jalba, A., 2011. Voxel-based assessment of printability of 3D shapes. In: *10th International Conference on Mathematical Morphology and Its Applications to Image and Signal Processing*, pp. 393–404.
- Umetani, N., Schmidt, R., 2013. Cross-sectional structural analysis for 3D printing optimization. In: *SIGGRAPH Asia Technical Brief*, article 5.
- Vanek, J., Galicia, J.A.G., Benes, B., Měch, R., Carr, N., Stava, O., Miller, G.S., 2014. PackMerger: a 3D print volume optimizer. *Comput. Graph. Forum* 33 (6), 322–332.
- Wang, W., Wang, T.Y., Yang, Z., Liu, L., Tong, X., Tong, W., Deng, J., Chen, F., Liu, X., 2013. Cost-effective printing of 3D objects with skin-frame structures. *ACM Trans. Graph. (SIGGRAPH Asia)* 32 (6), article 177.
- Xin, S.-Q., Lai, C.-F., Fu, C.-W., Wong, T.-T., He, Y., Cohen-Or, D., 2011. Making burr puzzles from 3D models. *ACM Trans. Graph. (SIGGRAPH)* 30 (4), article 97.
- Zhou, Q., Panetta, J., Zorin, D., 2013. Worst-case structural analysis. *ACM Trans. Graph. (SIGGRAPH)* 32 (4), article 137.
- Zhou, Y., Sueda, S., Matusik, W., Shamir, A., 2014. Boxelization: folding 3D objects into boxes. *ACM Trans. Graph. (SIGGRAPH)* 33 (4), article 71.