Computers & Graphics 58 (2016) 118–127

Contents lists available at ScienceDirect

Computers & Graphics

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

Special Issue on SMI 2016

Temporally coherent sculpture of composite objects

Artur P. Sampaio [a,b,*], Raphaëlle Chaine [a], Creto A. Vidal [b], Joaquim B. Cavalcante-Neto [b]

[a] Université Lyon 1, CNRS, LIRIS, UMR5205, F-69622, France
[b] Department of Computing, Universidade Federal do Ceará, 60455-760 Fortaleza, Brazil

ARTICLE INFO

Article history:
Received 15 March 2016
Received in revised form
9 May 2016
Accepted 10 May 2016
Available online 21 May 2016

Keywords:
Sculpting
Interactive techniques
Agglomerates

ABSTRACT

We address the problem of virtual sculpting and deformation of shapes composed of small, randomly placed objects. Objects may be tightly packed – such as pebbles, pills, seeds and grains, or be sparsely distributed with an overarching shape – such as flocks of birds or schools of fish. Virtual sculpture has rapidly become a standard in the entertainment industry. Composites, though, are still usually created in a static way by individually placing each object or by sculpting a support surface and procedurally populating the final shape. That raises problems for the generalization to evolving shapes with visual continuity of the components. Large amounts of geometrical data are generated, and must be maintained and processed, both by the CPU and by the GPU. Whenever the shape is deformed, one has to define how these compositing objects should turn, displace or disappear inside the volume, as well as how new instances should become visible to the outside. It is difficult to rely on a physical system to perform that task in real time. The system we suggest can be constructed upon any uniform mesh-based representation that can be deformed and whose connectivity can be updated by operations such as edge splits, collapses, and flips. The mesh remains populated with an aperiodic distribution of composing elements that are automatically updated under deformation. The idea is to sculpt the shape as if it were filled with little objects, without handling the complexity of manipulating volumetric shapes. For this purpose, we suggest exploiting the properties of the uniform sampling of the surface. We show that we are able to properly handle virtual sculpting of composites in real-time and maintaining temporal continuity. This system also uses GPU optimizations to render individual elements efficiently. To our knowledge, no previous sculpting system allows the user to simultaneously see and sculpt agglomerates in such a fast and reliable fashion.

1. Introduction

Modeling virtual objects is essential in the entertainment industry. Far from computer-aided design (CAD) and tedious static parameterization systems, artists expect to use abilities developed from real life sculpting in digital environments. Many systems now allow users to sculpt objects as they would do with virtual clay, but little effort has been made in the direction of materials composed of smaller entities that are glued together or that roll around each other. Artists should feel free to deform such global composites, as the elements change position, and appear or disappear automatically around each other (Fig. 1). Likewise, we would like to give artists the opportunity to sculpt a cloud of little objects that are not joined to each other or to deform such composites through other means, such as scripted deformations.

In this paper, we consider the sculpting and rendering of surfaces represented as agglomerates of smaller, composing 3D elements. These agglomerate materials are common in nature: rock piles, armies of ants, castles of sand, beautiful works of tiled art – basically any group of similar, randomly oriented elements that, when viewed together, form a larger scale shape (see Fig. 2). The 3D elements we use, henceforth denoted CompEls, can be seen as 3D models anchored at the faces of a support mesh. Representing and updating those kinds of assemblies efficiently remains an open problem. The sheer quantity of elements on a surface and the polygons required to represent them can cripple even powerful GPUs.

As stated in [3], the biggest challenge in creating an agglomerate is controlling the position of individual objects. Traditionally, agglomerates were created by placing each composing element in its final position, one by one. This can be cumbersome due to sheer object count. The other common approach is to design an intermediate representation and to populate it afterwards with smaller elements, using for example systems of particles in the case of a regular distribution or sample-based distributions [4]. This *a*

---

[*] Corresponding author. Tel.: +55 85 99139 8341.
E-mail address: artur.sampaio@liris.cnrs.fr (A.P. Sampaio).

http://dx.doi.org/10.1016/j.cag.2016.05.011
0097-8493/

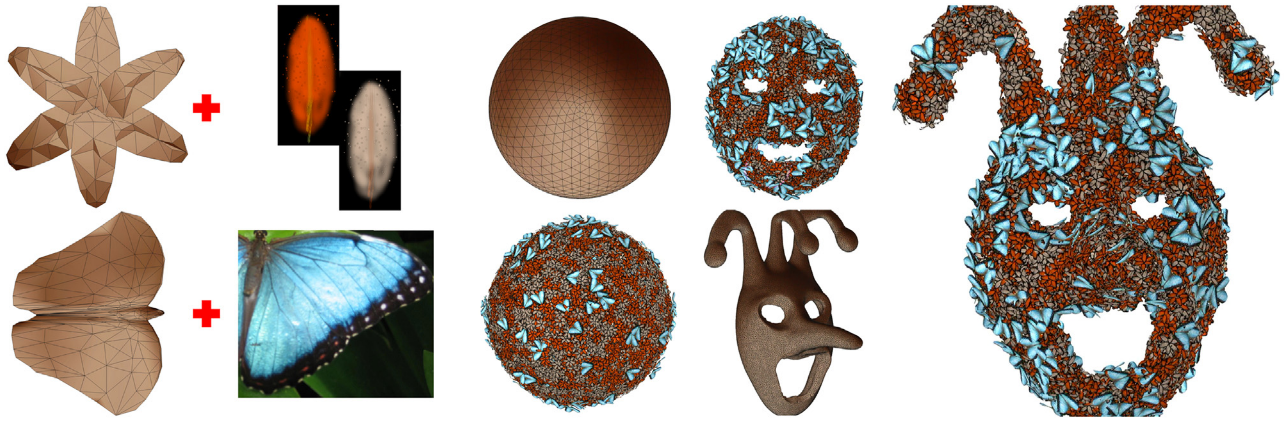

Fig. 1. Carnival mask sculpted using our framework. Left: input CompEls and accompanying textures. Center: initial sphere mesh and intermediate steps of the sculpting process, with and without CompEls on the outer surface. Right: final rendering.

An ant bridge reveals nature sculpting through agglomerates [1].

An ant bridge sculpted in our framework.

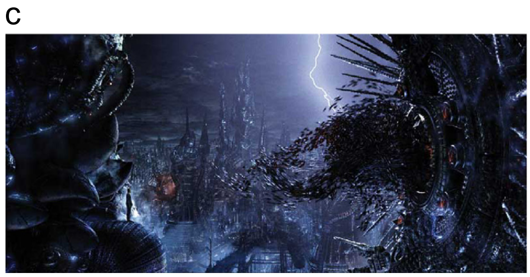

Deus Ex Machina [2]. The swarm of bugs forming the god's face illustrates an industry inspiration for our work.

Fig. 2. Agglomerates [1,2].

posteriori element sampling can either be limited to the object's surface or encompass the entire volume [5]. Extensive research has been done on how to generate and render such models correctly. The representation of the proxy object generally does not need to be exact or high quality. It can result from the rough meshing of an implicit surface, or from any other meshing design that does not involve constraints on the quality of the facets. The final rendered composite will display very high quality results and sampling properties. Those optimization techniques are clearly not intended for real-time applications or for dynamic surfaces. The corollary is that the user cannot see the desired agglomerate pattern before the end of the design process of the proxy shape. This can be a problem for the design of object clouds, where it becomes very difficult to have an idea of the final shape's outer silhouette without the presence of the actual objects. The problem is compounded on the representation of animations, where temporal continuity or coherence is also of concern. Merely resampling elements on updated regions of an animated surface can cause significant visual artifacts, due to sudden changes associated with rapid object insertion and removal, as described in [6]. The temporal continuity that we aim to obtain is similar to the one defined by Medeiros et al. [7], which means that objects should be inserted, displaced and removed without popping effects or loss of visual continuity between frames.

Another approach to minimize the cost of positioning the elements over the surface is to delay their involvement until the rendering pipeline, using works on spatial adaptive sampling of density function and projections. A number of publications focus on the process of rendering images using dots, called stipples. A stipple can later be transformed into the anchor of an orientable texture unit. For that reason, a good stippling can be the basis for agglomerate generation. One usually aims at obtaining blue noise samplings [8]. Hierarchical tiles with different levels of resolution [9,10] can be used to generate adaptive samplings with fast timings (except for preprocessing costs). More focused on meshed surfaces to be sampled, Pastor et al. [11] use a dynamic stippling generation system that performs directly on a mesh and that relies on adaptive subdivision meshes. However, although they use a hierarchical framework and the mesh could be edited for nearly isometric deformation, their work relies on fixed topology and overall connectivity. Finally, there is a tight coupling between vertices and stipples. For the same kind of applications, a hierarchical Poisson disk sampling was precomputed on the surface [7]. It can be later evaluated and rendered according to an importance function. Their system, however, uses fixed connectivity meshes. Abdrashitov et al. [12] employ a sketch-based interface to interactively create 2D mosaics. More recently, research has been done on detail-preserving deformation [13] that does not drastically affect the overall aspect of a shape, like the stretching of one part. The method is capable of maintaining