

Overview of Recent Work on Geometry Images

Hugues Hoppe
Microsoft Research

Extended Abstract

Surfaces in computer graphics are commonly represented using irregular meshes. While such meshes can approximate a given shape using few vertices, their irregularity comes at a price, since most mesh operations require random memory accesses through vertex indices and texture coordinates. Also, filter kernels must handle arbitrary mesh neighborhoods, and techniques like morphing, level-of-detail (LOD) control, and compression are complicated.

As an alternative, we have introduced the *geometry image* representation, which captures shape using a completely regular sampling, i.e. a 2D grid of (x,y,z) values [2]. The benefits of uniform grids are often taken for granted. Grids allow efficient traversal, random access, convolution, composition, down-sampling, and compression.

We briefly review our recent work on constructing geometry images from given input models, and highlight some benefits of this new regular shape representation.

The construction of a geometry image involves parametrizing a given surface onto a planar domain, and resampling the surface geometry on a regular domain grid. We have explored several such schemes.

Our original work [2] heuristically cuts an arbitrary surface into a disk using a network of cut paths, with $2g$ loops for a genus g surface. The resulting cut surface is mapped onto a square using a stretch-minimizing parametrization to reduce undersampling.

For shapes with high genus or long extremities, forcing the whole surface to map into a square can introduce high distortion. To mitigate this, we can instead cut the surface into several pieces to produce a *multi-chart* geometry image [6]. The challenge is to join these piecewise regular charts into a watertight surface.

For genus-zero models, a geometry image may be constructed via *spherical parametrization* [5], which does not require any a priori surface cuts. The spherical domain is unfolded into a square using a simple cut with elegant boundary symmetries. These boundary symmetries permit the construction of a smooth (C^1) polynomial surface, and the regular control grid structure lets the surface be evaluated entirely within the GPU rasterization pipeline [4]. In addition, a spherical geometry image can be compressed using traditional image wavelets, to produce a concise shape description [3].

Geometry images for static objects can be generalized to *geometry videos* for animated shapes [1]. Inter-frame prediction significantly helps the compression of such animations.

In conclusion, as the cost of 3D transformations becomes negligible in emerging hardware, one should re-evaluate whether geometry would not be better represented using ordinary grids. Rendering from such geometry images involves simple sequential data traversal, without index dereferencing and texture coordinate lookups. The regular grid structure also simplifies LOD, compression, and morphing. More generally, it enables the powerful GPU rasterization architecture to process geometry in addition to images..

(See also <http://research.microsoft.com/~hoppe/>.)

References

- [1] BRICEÑO, H., SANDER, P., McMILLAN, L., GORTLER, S., AND HOPPE, H. 2003. Geometry videos: A new representation for 3D animations. *ACM Symposium on Computer Animation 2003*, pp. 136-146.
- [2] GU, X., GORTLER, S., AND HOPPE, H. 2002. Geometry images. *ACM SIGGRAPH 2002*, pp. 355-361.
- [3] HOPPE, H., AND PRAUN, E. 2003. Shape compression using spherical geometry images. *Multiresolution in Geometric Modeling 2003*, to appear.
- [4] LOSASSO, F., HOPPE, H., SCHAEFER, S., AND WARREN, J. 2003. Smooth geometry images. *Eurographics Symposium on Geometry Processing 2003*, pp. 138-145.
- [5] PRAUN, E., AND HOPPE, H. 2003. Spherical parametrization and remeshing. *ACM SIGGRAPH 2003*, pp. 340-349.
- [6] SANDER, P., WOOD, Z., GORTLER, S., SNYDER, J., AND HOPPE, H. 2003. Multi-chart geometry images. *Eurographics Symposium on Geometry Processing 2003*, pp. 146-155.