

NRC Publications Archive Archives des publications du CNRC

A Fast Hybrid Geomorphing LOD Scheme

Borgeat, Louis; Fortin, P.-A.; Godin, Guy

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. /
La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version
acceptée du manuscrit ou la version de l'éditeur.

Publisher's version / Version de l'éditeur:

SIGGRAPH 2003 Sketches and Applications [Proceedings], 2003

NRC Publications Archive Record / Notice des Archives des publications du CNRC :

<https://nrc-publications.canada.ca/eng/view/object/?id=88cc5139-bdbd-4a29-904a-bd85e167424c>

<https://publications-cnrc.canada.ca/fra/voir/objet/?id=88cc5139-bdbd-4a29-904a-bd85e167424c>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at

<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site

<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at

PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

A fast hybrid geomorphing LOD scheme

Louis Borgeat*
National Research Council of Canada

Pierre-Alexandre Fortin
Laval University, Canada

Guy Godin
National Research Council of Canada

We present a new technique for fast, view-dependent, real-time visualization of large multiresolution geometric models with color or texture information. This method combines the rendering speed of optimized static models, the smooth transitions associated with geomorphing, and the viewpoint adaptivity of iterative vertex contraction based methods.

1 System overview

A significant amount of recent work on real-time multiresolution rendering of triangulated geometric models is built on the concept of iterative edge contraction (e.g. [Hoppe 1997]). The system we have developed is a hybrid method that integrates aspects of this kind of approach and of more classical discrete Level of Detail (LOD) techniques. Our goal was to produce a multiresolution rendering system that would produce minimal visual artifacts when rendering high resolution scene or object datasets produced from sensor data while maintaining the same graphic performance we get with optimized static versions of those models.

The first step of the method is to partition the triangular mesh into a set of groups. The model is then globally decimated into a series of discrete LODs using an algorithm based on vertex pair contraction. Each discrete level is repartitioned along the same boundaries. Groups are shaped based on criteria such as compactness, orientation, texture, and desired granularity.

At run-time, LOD levels and geomorph ratios between selected levels are computed for each group of the model, these groups are then rendered as precomputed triangle strips. Border points between groups are geomorphed in order to maintain seamless continuity between neighboring LOD groups at all time. To achieve this, we maintain a connexity graph between the various groups and apply a set of simple rules when morphing these points. Morphing is done for space and texture coordinates, normals and color.

Adjusting the granularity of the multiresolution rendering through this grouping process provides us with multiple advantages. First, it allows us to extend smooth geomorphs over numerous frames by morphing more triangles over a longer distance rather than expanding/contracting the model by a few polygons every frame. Secondly, each group having a static topology, it can be efficiently rendered as a set of pre-computed triangle strips. These groups can also serve as the units of an anticipative paging algorithm for larger models. Finally, groups can be shaped so that they correspond to rectangular texture images, therefore optimizing texture memory usage and smoothing texture swapping for high resolution models.

2 Results

Our current implementation uses quadric error metrics [Garland and Heckbert 1998] to decimate the models and SGI's OpenGL Performer as the base of the run-time infrastructure. The multiresolution processing is done asynchronously on a different CPU than the one used for rendering. We tested our method on various colored or textured models produced using range sensors and digital

photographs. Our test system was a 1.4 GHz dual Athlon PC with a GeForce 4 graphic card. The figure shows a slanted view of a rectangular coneiform tablet to illustrate the similar pixel size of the triangles at different distances from the observer. The model is textured and varies in size from 25k to 2.5M polygons. The tablet, scanned using the National Research Council's color range sensor, was provided by Prof. Marc Levoy of Stanford University while it was on loan from Renee Kovacs.



During those tests, the time required to perform the LOD selection and the morphing of all visible triangles (geometry, color/texture and normals) on one of the CPUs amounted to roughly 30% of the time required to render the corresponding frame, allowing free cycles for other processing. The number of rendered polygons per second was the same for a static model loaded with Performer's viewer and for its multiresolution counterpart loaded with our application. This confirms that there is no graphic penalty associated with the multiresolution processing, therefore allowing a maximal frame rate increase. Also, as expected, this method produces almost unnoticeable transformations during LOD morphing, and produces especially good results with textured data.

We are currently implementing the anticipative paging and compression strategy that will allow navigation through larger models. We are also pursuing work on automatic segmentation of the models, on recursive subdivision of the groups, and we are implementing more advanced LOD selection criteria based on both resolution and frame rate constraints.

References

- GARLAND, M., AND HECKBERT, P. S. 1998. Simplifying surfaces with color and texture using quadric error metrics. In *Proceedings of the conference on Visualization '98*, IEEE Computer Society Press, 263–269.
- HOPPE, H. 1997. View-dependent refinement of progressive meshes. *Computer Graphics 31*, Annual Conference Series, 189–198.

*e-mail: louis.borgeat@nrc-cnrc.gc.ca