Virtual Backdrops

Daniel G. Aliaga, Anselmo A. Lastra

Department of Computer Science
University of North Carolina
Chapel Hill, NC 27599-3175
{aliaga | lastra } @cs.unc.edu

Large and complex 3D models are required for architectural walk-throughs, flight simulators and other applications of virtual environments. It is not possible to render all the geometry of these arbitrarily complex scenes at highly interactive rates, even with high-end computer graphics systems. This has led to extensive work in 3D model simplification methods.

We have been investigating dynamically replacing portions of 3D models with textures. This approach is similar to the old stage-trick of draping cloth backgrounds in order to generate the illusion of presence in a scene too complex to actually construct on stage. A texture (or backdrop) once created, can be rendered in time independent of the geometric complexity it represents. Consequently, significant frame rate increases can be achieved at the expense of memory and image quality. Properly handling these last two tradeoffs is crucial to a successful texture-based simplification algorithm.

Previous work has approached these tradeoffs by requiring a large number of texture samples in order to guarantee a maximum per-pixel error. While for some ranges of applications and models this provides a good solution, we have found that in our domain (architectural models, CAD models and other indoor models) the simplification benefits are overwhelmed by the large number of textures that need to be precomputed or dynamically rendered.

A successful algorithm for dynamically replacing geometry with textures must provide solutions for the following three major problems:

- Geometric Continuity: a texture contains an image of a portion of the model. The image is only perspectively correct when viewed from the same viewpoint used to create it. Thus, when the eye moves from the texture's original viewpoint, the geometry adjacent to the texture appears discontinuous with the image of the texture. Previous systems addressed this issue by either allowing a small discontinuity or restricting the texture locations to fully contain entire "objects".

- Temporal Continuity: when the eye approaches a texture (or recedes far enough from a subset of the model), we need to switch the texture to geometry (or vice versa). During this motion, the viewpoint will not be at the same position from where the texture was created. This sudden transition will cause an effect commonly known as "popping".

- Automatic Placement: the texture-based simplification system will need to choose where in the model to place textures in order to achieve the desired speedup and image quality.

Solutions have been proposed in [Aliaga96] for the Geometric and Temporal Continuity problems (see Figure 1). We have been focusing on the automatic placement aspect. For our purposes, we can loosely classify indoor models into two categories:

- Single-Room Models: models where most of the geometric complexity is visible from the typical viewing locations (few large occluders are present).

- Multiple-Room Models: this category fits well with the cells and portals framework used by some visibility culling algorithms.

The special characteristics of the cells and portals framework where used to solve the automatic placement problem for cell-partitioned models (see Figure 2) [Aliaga97]. We are currently exploring solutions for the texture placement problem in general 3D models.


Figure 1: Geometric Continuity. Two textures outlined in red. Adjacent geometry warped to match the textures.

Figure 2: Portal Textures. The portals have been replaced with textures.

[Aliaga96] D. Aliaga, "Visualization of Complex Models Using Dynamic Texture-based Simplification", *IEEE Visualization '96*, pp. 101-106, 1996.

[Aliaga97] D. Aliaga, A. Lastra, "Visibility Culling using Portal Textures", UNC TR# 97-009, April 1997, submitted for publication.