



Optimized GPU evaluation of arbitrary degree NURBS curves and surfaces

Adarsh Krishnamurthy*, Rahul Khardekar, Sara McMains

Computer Aided Design and Manufacturing Lab, University of California, Berkeley, United States

ARTICLE INFO

Article history:

Received 16 July 2008

Accepted 13 June 2009

Keywords:

NURBS

GPU

Surface evaluation

Level of detail

ABSTRACT

This paper presents a new unified and optimized method for evaluating and displaying trimmed NURBS surfaces using the Graphics Processing Unit (GPU). Trimmed NURBS surfaces, the de facto standard in commercial mechanical CAD modeling packages, are currently being tessellated into triangles before being sent to the graphics card for display since there is no native hardware support for NURBS. Other GPU-based NURBS evaluation and display methods either approximated the NURBS patches with lower degree patches or relied on specific hard-coded programs for evaluating NURBS surfaces of different degrees. Our method uses a unified GPU fragment program to evaluate the surface point coordinates of any arbitrary degree NURBS patch directly, from the control points and knot vectors stored as textures in graphics memory. This evaluated surface is trimmed during display using a dynamically generated trim-texture calculated via alpha blending. The display also incorporates dynamic Level of Detail (LOD) for real-time interaction at different resolutions of the NURBS surfaces. Different data representations and access patterns are compared for efficiency and the optimized evaluation method is chosen. Our GPU evaluation and rendering speeds are more than 40 times faster than evaluation using the CPU.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

Non-Uniform Rational B-Splines (NURBS) are the industry standard for the representation of geometry in mechanical Computer Aided Design (CAD) systems. Although NURBS are ubiquitous in the CAD industry, there is currently no built-in hardware support for displaying NURBS surfaces. OpenGL provides a software NURBS solution; however, the implementation is not fast enough for evaluating large surfaces interactively, and in our experience it often renders trimmed NURBS surfaces incorrectly. Because surface evaluation is a computationally intensive operation, the common practice in CAD systems is to preprocess the NURBS surfaces by evaluating and tessellating them into triangles, and then using the standard graphics pipeline to display them.

The use of a preprocessing technique not only leads to very high memory usage, but also restricts the surface evaluation to a particular Level of Detail (LOD). Hence, a highly enlarged view of the surface may not be tessellated sufficiently, whereas a distant view may render an excessive number of triangles. In this paper, we describe a method by which we evaluate and display a trimmed NURBS surface directly, without approximating it by simpler surfaces, using a programmable graphics card. The usage of the GPU's computational power not only speeds up the surface

evaluation significantly but also reduces the CPU memory usage, eliminating the need for calculating and storing the tessellation data or simplified surface information that is typically used only for visualization purposes.

Previous GPU methods [1,2] focused mainly on rendering NURBS surfaces rather than exact evaluation. Hence, they approximated a higher degree NURBS surface by lower degree Bezier surfaces that closely resemble the original surface based on pixel location error while rendering. Even though such approximations are good enough for rendering, they cannot be extended to a general-purpose NURBS evaluator capable of handling arbitrary degree NURBS surfaces. We introduced a unified method to evaluate arbitrary degree NURBS surfaces on the GPU without making any approximations [3]. The contemporaneous work by Kanai [4] for evaluating NURBS surfaces also did not use any approximations, but required different GPU programs for evaluating NURBS surfaces of different degrees. This makes the implementation of their system tedious, since specific new programs have to be written for surfaces of different degrees. Moreover, since standard CAD models can be made of surfaces of widely varying degrees, with surfaces up to degree 100 occurring in many complex models, a unified NURBS evaluation algorithm will be a more practical solution.

In this paper we describe our unified NURBS evaluation and rendering method, expanded from the original conference presentation [3]. The main contributions of our approach include:

- A GPU method for evaluating arbitrary degree NURBS surfaces with an arbitrary number of control points and knots with

* Corresponding author. Tel.: +1 510 590 7325.

E-mail addresses: adarsh@me.berkeley.edu (A. Krishnamurthy), rahul@me.berkeley.edu (R. Khardekar), mcmains@me.berkeley.edu (S. McMains).

the same unified fragment program. Our method uses the GPU to evaluate a grid of points on the NURBS surface that can be directly used for rendering as well as for further modeling operations. Our method is easily extensible to evaluate derivatives and normals of the NURBS surface.

- Backward-compatible algorithms that make use of standard OpenGL extensions or features that are available even in cards that are more than 5 years old, while still taking advantage of the improved performance on newer cards.
- Different implementations of the evaluation algorithm that use different memory access patterns and data packing on the GPU. We choose the optimum evaluation method based on the performance of these different implementations.
- A direct method to render trimmed NURBS surfaces by interpreting the points already evaluated as vertices. The rendering algorithm is capable of dynamic continuous LOD based on the size and location of the surface with respect to the view point.

2. Background and related work

2.1. Programmable GPUs

Graphics processing units (GPUs) have recently evolved into programmable parallel processors capable of performing general-purpose computational tasks [5,6]. We make use of two programmable units on the GPU, the Vertex Processing Unit (VPU) and the Fragment Processing Unit (FPU), which can execute a user-defined set of instructions, called the vertex program and the fragment program, for each vertex and fragment respectively, in the place of a fixed sequence of geometric transformations, lighting operations (per-vertex operations), and texturing operations (per-fragment operations). Vertex programs can obtain the geometry and attribute (color, texture coordinates, etc.) data stored in the GPU memory via traditional display lists or more recently, Vertex Buffer Objects (VBOs). Geometric primitives (triangles generally) assembled from the vertex data then get rasterized into fragments (potential pixels) that pass through the FPU. Vertex and fragment programs can access data stored in textures that can have full 32-bit floating point precision. Usually the output of the FPU goes into a frame buffer, which is a 2D block of memory with four attributes at each location. In modern GPUs, the FPU can also output directly to a floating point texture (render-to-texture) using off-screen render targets called Frame Buffer Objects (FBOs). This allows the use of the output of a first pass through the rendering pipeline as input texture data for the second pass. FBOs can also be used to render into a Vertex Buffer Object (VBO) so that the output can be used as vertex data for the next rendering pass. Because multiple vertices and pixels are processed in parallel, and operands are four-component vectors, GPUs can achieve much higher computational speeds than conventional CPUs on arithmetically intensive operations.

2.2. NURBS evaluation techniques

Many early high-quality renderings of curved surfaces used ray tracing. Toth [7] and Nishita et al. [8] perform ray tracing on parametric and rational surfaces by solving for the ray-surface intersection point using numerical methods. Martin et al. [9] gives a complete algorithm for ray tracing trimmed NURBS. Pabst et al. [10] used ray casting on the GPU to render trimmed NURBS surfaces.

To take advantage of graphics hardware, parametric surfaces tend to be tessellated before display. Much work on trimmed NURBS focuses on the trimming aspect. The OpenGL version 1.1 implementation renders trimmed NURBS surfaces using the method presented by Rockwood et al. [11] for trimmed parametric surfaces, which divides the parametric domain into

patches based on the trim curves. These patches are then tessellated in the 2D domain and then evaluated to find the surface point coordinates. However, in our experience the OpenGL implementation tessellates trimmed NURBS surfaces incorrectly at trim curve concavities. In addition, being a CPU evaluator, it is not fast enough to render large numbers of trimmed NURBS surfaces at interactive rates.

Previous work such as [12–14] displayed NURBS after first converting them to Bezier patches and converting the trimming curves to Bezier segments, since Bezier evaluation is less computationally demanding. These patches were then triangulated and sent to the graphics card for display. Guthe et al. [1,2] approximate each NURBS surface with lower degree Bezier patches, but they then evaluate the Bezier patches on the GPU after the CPU approximation step. They also introduced a LOD system for choosing the appropriate approximation patch decomposition and the sampling density. Since in general no Bezier surface of lower degree can exactly match an arbitrary degree NURBS surface, a disadvantage of this approach is that the final surface may not achieve sufficient accuracy unless it is split into many Bezier patches, increasing the number of patches by up to two orders of magnitude in their examples.

Subdivision surfaces, which have largely replaced tensor-product patches in entertainment applications where mathematical exactness is not required, have also been directly evaluated on the GPU. Prior work by Bolz and Schröder [15] and Shiue et al. [16] focused on using a fragment program to compute the limit points of Catmull–Clark subdivision meshes. These methods can be extended to evaluate uniform B-spline surfaces; the limit surface of a Catmull–Clark subdivision in the absence of extraordinary points is the bi-cubic B-spline surface. However, they cannot be extended to evaluate NURBS because they do not have a subdivision scheme with stationary rules [17,18]. Loop and Blinn [19] used the GPU to render piecewise algebraic surfaces of lower degrees. However, it is difficult to extend the method to evaluate arbitrary degree NURBS surfaces.

The fragment-program implementations of surface evaluation of subdivisions were not fast enough for real-time interaction with a large number of surfaces because the evaluated surface coordinates had to be read back from an off-screen pixel buffer using an expensive p-buffer switch for each surface. Guthe et al. [1] overcome this issue by using a vertex program, but their method is not as flexible because the number of parameters that can be passed to a vertex program is quite limited, and vertex texture fetches are possible only in the latest graphic cards. Thus, they approximated the original input by a hierarchy of bi-cubic Bezier patches to limit the amount of data that needed to be transferred per patch. In our approach, we use a fragment program but get around the p-buffer switch issue by using a frame buffer object, which renders directly to a texture, and a vertex buffer object, which takes this texture as input coordinates for a subsequent rendering pass.

Recently, Kanai [4] developed a fragment-program based NURBS evaluation that closely resembles our method. However, their implementation required different fragment programs for surfaces of different degrees. While this method is theoretically capable of evaluating any NURBS surface, its implementation becomes tedious since different fragment programs have to be written specifically for each possible degree of a NURBS surface that may be present in a model. Hence a unified evaluation method that can be used to evaluate arbitrary degree NURBS surfaces is preferred.

2.3. NURBS curve and surface definitions

In this section, we briefly review the mathematical notation used for defining NURBS curves and surfaces, adapted from Piegl

Download English Version:

<https://daneshyari.com/en/article/440371>

Download Persian Version:

<https://daneshyari.com/article/440371>

[Daneshyari.com](https://daneshyari.com)