

Abstract

Feed-forward rendering through projection followed by rasterization is the approach of choice in interactive 3-D computer graphics. The most frequently used camera model is the planar pinhole, which is simple and which produces images familiar to the user. However, the requirement that all rays pass through one point is restrictive in many computer graphics contexts. For example, a pinhole camera cannot produce a robust image-based approximation of a 3-D scene that supports viewpoint translation without disocclusion errors, it cannot model a soft-shadow map, and it cannot model rays reflected off non-planar surfaces. To address these limitations, several non-pinhole camera models have been recently introduced. Although non-pinholes, these camera models do provide efficient projection, which opens the door to feed-forward rendering with graphics hardware support. In this paper we discuss the challenges of feed-forward non-pinhole rendering and we describe general approaches for overcoming them. We also describe sample implementations that take advantage of the programmability at vertex, primitive, and fragment level exposed by current graphics processing units.