

Paul Rosen
The Graph Camera

ABSTRACT

In interactive 3-D graphics applications the user typically explores the scene by positioning and orienting a virtual camera. When the experience of actual locomotion in the virtual space is unnecessary, such sequential exploration is undesirable since it is inefficient—the user has to cover large distances in the scene, and ineffective—the user can only see a small fraction of the scene at any given time, which is particularly inadequate for dynamic scenes. The conventional solution is to employ several stationary cameras that render the scene in parallel. However, a large number of cameras is required for adequate scene coverage and there is no continuity between individual images, which requires the user to adapt to a multitude of contexts, one at the time. We introduce the *graph camera*, a non-pinhole camera with rays that circumvent occluders to sample most or all of a 3-D scene. The graph camera image has a single layer, it is mostly continuous and non-redundant, yet it shows simultaneously all regions of interest in a complex 3-D scene. The graph camera is constructed from a planar pinhole camera through a series of view frustum bending, splitting, and merging operations. The graph camera has tens or even hundreds of frusta, yet rendering is efficient due to a fast projection operation that allows rendering in a single pass and allows resolving visibility automatically.