Projector-Based Dual-Resolution Stereoscopic Display

Guy Godin[*] Jean-François Lalonde Louis Borgeat
Visual Information Technology Group
National Research Council of Canada
Ottawa, Ontario, Canada

Abstract

We present a stereoscopic display system which incorporates a high-resolution inset image, or fovea. We describe the specific problem of false depth cues along the boundaries of the inset image, and propose a solution in which the boundaries of the inset image are dynamically adapted as a function of the geometry of the scene. This method produces comfortable stereoscopic viewing at a low additional computational cost. The four projectors need only be approximately aligned: a single drawing pass is required, regardless of projector alignment, since the warping is applied as part of the 3-D rendering process.

1. Introduction

The size and complexity of models displayed in virtual environments have increased dramatically, in part due to the emergence of 3-D digitizing tools and techniques which efficiently produce detailed models of real scenes and objects. However, this growth has not been matched by a similar increase in the resolution of currently available displays. The lowering cost of commodity projectors and graphic hardware has encouraged the construction of large-scale mosaic walls as one solution. Alternately, the apparent resolution of a display can be locally increased by inserting a high-resolution inset within a larger, lower-resolution image. This approach has been used in flight simulators and head-mounted displays (e.g. [3][4]). Recently, similar projector-based systems have been proposed, for example "focus plus context" screens [2] and "foveal displays" [1].

This paper extends this class of methods by introducing a dual-resolution *stereoscopic* display system (Fig. 1). However, this configuration poses a specific new problem: if a high-resolution inset is added in each eye's view, ambiguous depth perception may arise along the boundary with the lower resolution image and interfere with scene perception. We discuss this issue, and outline the general and computationally efficient solution that we have developed.

[*]Correspondence to: Guy.Godin@nrc-cnrc.gc.ca

Figure 1. Dual-resolution stereoscopic display with adaptive boundaries (screen photograph).

2. Dual-resolution stereo display

A dual-resolution polarization-based passive stereoscopic display requires four projectors arranged such that a high-resolution small inset image is projected within a larger image (Fig. 2). This inset is often referred to as a *fovea*, in analogy to the biological visual system. The four images are generated on separate synchronized 3-D graphic computers. The projectors may be only approximately positioned and aligned, the warping and trimming required to align the images on the screen (left vs right and low- vs high-resolution) are estimated using camera-based techniques. Our system applies the image warping at no additional cost, by modifying the projection matrices used for the 3-D rendering pass. Full consistency of the coordinate systems is obtained by enforcing identical near and far clipping planes.

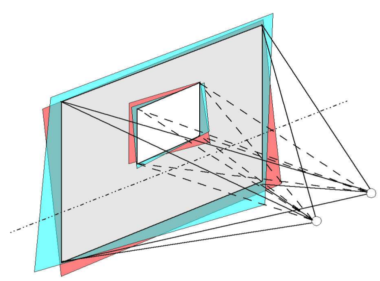

Figure 2. Projection geometry.

Because of its smaller area, the inset is significantly brighter than the larger image. Perfectly matching the brightness and colour of the four projectors is very difficult;

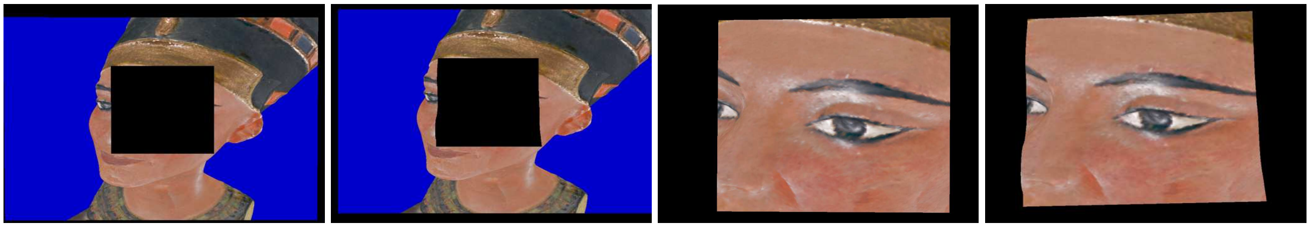

Figure 4. The four images forming the view appearing in Fig. 1.

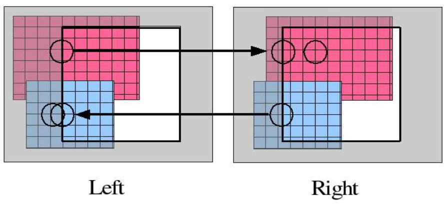

Figure 3. Matching ambiguity at the boundary.

we choose to preserve this additional brightness, useful in enhancing the area of interest and for lights-on operations.

However, the brightness difference between the inset and the periphery creates perceptible edge features in both left and right views. These rectangular contours become strong stereoscopic cues that interfere with the underlying stereoscopic image of the scene being displayed: for a given screen location, two different matches are possible, based either on the intensity edge or on the scene contents (Fig. 3). This ambiguity creates two competing layers of perceived depth along the boundary of the inset image. Applying large smooth cross-fading zones only slightly attenuates the problem. Even in the case when the brightness of all four projected images are similar, the stereo fusion of high and low resolution images does not provide the same perceived level of details as the fusion of two high-resolution images, and the boundary may still be noticeable.

In both cases, the problem arises from displaying an element of the scene in the low-resolution image for one eye and in the (brighter) high-resolution image for the other. This statement of the problem hints at its solution: the boundaries of the high-resolution insets must be displaced so that they are located over the same scene positions in both views. The boundary can be virtually moved within the footprint of the inset projector by inserting a corresponding black area. This solution requires that the screen position of this boundary be recomputed whenever the underlying scene or the observer's position change, in the worst case for every displayed frame. Finding a computationally efficient method for adjusting this boundary is therefore essential.

3. Proposed method

The method is implemented as a post-rendering pass. Each of the four images is rendered for the corresponding eye location and field of view, with the warping required for on-screen alignment. To avoid accessing directly the scene geometry, we use the depth buffer of the current image as a proxy. The original 3-D coordinates of an image pixel can be obtained by applying the inverse of the projection matrices. Matching points along epipolar lines are found for the left and right views. In absence of occlusion, the sought boundary points are found in constant time, whereas occlusion may require a linear search. Only a subset of the depth buffer needs to be read back.

In the superimposed left and right views (Fig. 1), the shape adaptation of the inset boundaries is visible, especially along the cheek. The four source images forming this display are shown in Fig. 4, which also illustrates the effects of the warping and the black masking on individually computed views. The scene-adaptive displacement of the two vertical boundaries appears clearly. A small amount of feathering can be applied around the adapted boundaries to smooth the transition, but for clarity it is omitted here.

A stereoscopic dual-resolution system allows enhanced stereoscopic viewing of detailed 3-D models in the inset while retaining the overall view in a unified stereoscopic display. However, fixed inset boundaries cause discomfort and severely limit the usefulness of the approach. This problem is solved by our proposed adaptive modification of the boundary, which is an *essential* component of a stereoscopic inset. The general solution proposed here adds only a small computational cost, even for complex models.

References

[1] M. Ashdown and P. Robinson. The Escritoire: a personal projected display. In *11th Int. Conf. in Central Europe on Computer Graphics, Visualization and Computer Vision*, pp. 33–40, 3-7 Feb. 2003.

[2] P. Baudisch, N. Good, and P. Stewart. Focus plus context screens: combining display technology with visualization techniques. In *Proc. of the 14th Annual ACM Symposium on User Interface Software and Technology*, pp. 31–40, Nov. 2001.

[3] Evans and Sutherland. Vistaview, image projection system for full-flight simulation, http://www.es.com/products/displays/vistaview/, Dec. 2003.

[4] A. Yoshida, J. Rolland, and J. Reif. Design and applications of a high-resolution insert head-mounted-display. In *Virtual Reality Annual International Symposium '95*, pp. 84–93, 1995.