Real-Virtual Antialiasing

Jan Fischer, University of Victoria, Canada

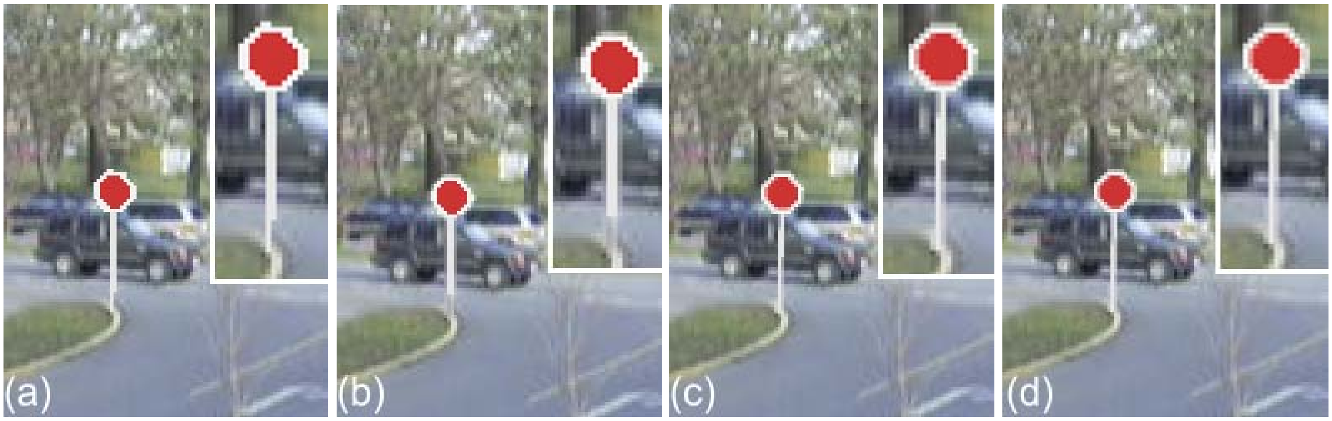

Figure 1: *Computer-generated stop sign rendered over video image. (a) No antialiasing. (b) Original resolution averaging. (c) 2x real-virtual supersampling. (d) 4x real-virtual supersampling. In the top right corner of each figure, a magnified image detail is shown.*

Abstract

The integration of virtual renderings into video images is of great importance for many applications of computer graphics such as TV and movie production and augmented reality. This submission proposes and examines methods for a smooth transition from real to virtual image regions along the boundaries of graphical objects. This constitutes a new form of antialiasing, which - unlike its counterpart in purely virtual images - has not been systematically investigated so far. We introduce real-virtual supersampling based on the combination of a high-resolution rendering of the graphical elements and an upsampled version of the video image. Several different methods for image upsampling are discussed. We compare results obtained through these methods with simple image combination methods: image overlay without antialiasing and pixel averaging in the original image resolution.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms

Keywords: antialiasing, video augmentation, augmented reality

1 Description

In many applications of computer graphics, computer-generated virtual objects are rendered over captured real images. The combination of real and virtual image elements is for instance essential in TV and movie production, as well as in mixed and augmented reality applications. In many cases, graphical models are superimposed over the real image with straightforward rendering methods. This leads to aliasing at the boundaries between real and virtual image regions, as illustrated in Fig. 1 (a). Unless image postprocessing filters are applied (e.g., superimposed atmospheric effects), this aliasing can diminish the quality of the visual integration of virtual objects into captured video sequences. This is particularly true if the capturing or display devices do not offer high image resolutions.

The work presented here deals with automatic real-time antialiasing at the boundary between real and virtual image regions. Real-virtual supersampling is enabled by upsampling the real video image and rendering a high resolution version of the virtual content over it. This results in a combined high resolution image, which is used as input for the final image composition step. In the image composition step, camera image and virtual object are rendered at the original, smaller resolution. Only for pixels at the boundary between real and virtual image regions, the high resolution image is accessed, and the pixel colors at the corresponding location are averaged. The averaged high resolution pixels are used as boundary pixels for the combined image, which leads to an antialiased output (see Fig. 1 (c),(d)).

The real-virtual antialiasing approach was implemented using image processing shaders on the GPU, resulting in real-time frame rates even for large image resolutions and upsampling factors. In order to compare the image quality generated with real-virtual antialiasing, two alternative image combination methods were also implemented. The standard approach (Fig. 1 (a)) renders virtual models over the camera image without antialiasing. A simple antialiasing approach (Fig. 1 (b)) performs image space averaging at object boundaries in the original image resolution. Original resolution averaging, however, blurs the borders and leads to disappearing small structures. If real-virtual supersampling is used, a better image quality was found with 4x supersampling (Fig. 1 (d)) compared to 2x supersampling (Fig. 1 (c)).

Several different image upscaling methods were implemented for real-virtual supersampling. Straightforward sample-and-hold (nearest neighbor) interpolation, biquadratic B-Spline interpolation, and a more complex adaptive edge-directed upscaling scheme were investigated. The GPU-based edge-directed upscaling method was implemented based on principles described in [Kraus et al. 2007]. Currently, a method for a quantitative evaluation of the comparative image quality of different antialiasing methods is being developed.

References

KRAUS, M., EISSELE, M., AND STRENGERT, M. 2007. GPU-Based Edge-Directed Image Interpolation. In *Proc. of SCIA*.