Volumetric and User-Centric Rendering Techniques for Lens Flare and Film Grain in Virtual Reality Environments

Johann Wentzel
http://johannwentzel.ca
School of Computer Science,
University of Waterloo

Lesley Istead
https://carleton.ca/hci/people/lesley-istead/
School of Information Technology,
Carleton University

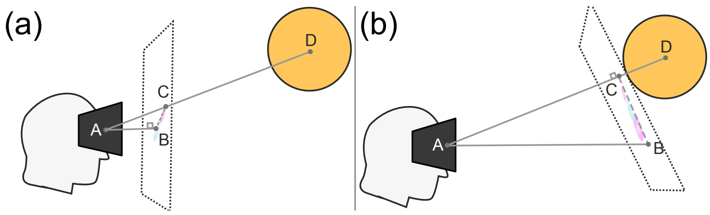

Figure 1: Two techniques for rendering lens flare in VR: (a) *Headset Flare*, which mounts the rendering plane in front of the user's head; and (b) *Directional Light Flare*, which mounts the rendering plane in front of the sun. The angle between the user's gaze direction (*AB*) and the direction toward the sun (*AD*) determine the spread and angle of the flare's rings, expanding outward from plane intersection point *C* along *CB*.

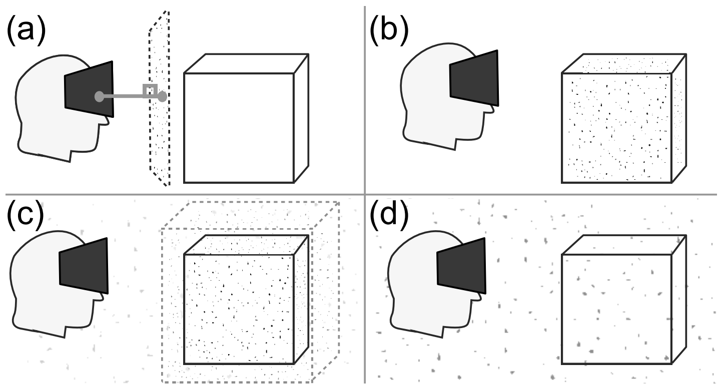

Figure 2: Four techniques for rendering film grain in VR: (a) *Headset Grain*; (b) *On-Surface Grain*; (c) *Object Cloud*; and (d) *Free Cloud*.

While several techniques can recreate lens flare and film grain effects in both monoscopic [1] and stereoscopic [2] formats, little work has explored applying these effects in real-time immersive environments like virtual reality (VR). Developing visual effects for VR involves several distinct challenges compared to other formats, the most prominent of which is user comfort. If a user moves in space, how should an effect react in order to remain visible, while ensuring that the user can fuse their left- and right-eye views? To address this, we present several methods for producing lens flare and film grain effects in VR. Lens flare is rendered in two different styles, taking as input the headset's forward direction and vector toward a light source. The primary difference between these styles is the position of the rendering plane, either mounted to the user's headset (*Headset Flare*) or mounted on the light source itself (*Directional Light Flare*). Film grain is rendered using two techniques: shader techniques which render grains as a 2D texture (*Headset Grain* and *On-Surface Grain*); and volumetric techniques which render grains as 3D particles within the environment (*Object Cloud* and *Free Cloud*).

We simulate lens flare with two artifacts rendered in the environment. The first is a series of semi-transparent rings with chromatic scattering similar to those found in analog lens flare. Ring positioning is based on three components (Figure 1): the vector of the user's headset gaze direction (*AB*), the vector between the user's headset and the sun (*AD*), and the intersections of those vectors with a rendering plane. Ring artifacts are distributed on the rendering plane the along the vector between these intersection points (*B* and *C*), with the size of $\angle CAB$ determining their spread. From the user's perspective, this results in ring artifacts splayed from the light source toward the center of their camera view, increasing in spread as the user moves their gaze away from the light. These rings are rendered in world space such that the user's eyes can fuse the two lens flare images. The second effect, a static glare artifact accompanied by a volumetric blur, is rendered at point *C* to appear to the user as surrounding the sun. Both effects fade in and out upon $\angle CAB$ reaching a set size, which in our implementation was 15 degrees.

We implemented lens flare in two styles, the difference between them being the position and anchoring of the rendering plane. *Headset Flare* (Figure 1a) displays these artifacts on a plane about 5 cm in front of the user's headset position, anchored to and centered with the user's headset view. From the user's perspective, flare effects would appear near their headset, close to their eyes. *Directional Light Flare* (Figure 1b) positions this plane near the sun, distant from the user. From the user's perspective, flare effects would appear distant, spread across the sky.

We simulate film grain using two rendering techniques. The *Shader* technique uses Perlin noise to procedurally generate grains of varying shape, brightness, and transparency. This shader's intensity and number of grains varies based on the distance from the camera to the object to which this shader is applied. This results in a flat texture containing small, semi-transparent grains simulating film grain. The *Volumetric* rendering technique creates small 3D grain objects, randomly scaled between 0.5 and 3 cm in all 3 dimensions, with random brightness and transparency. When applied to a volume, this technique creates a bounding volume V_b of a user-configurable scale larger than the original volume V_o (in our implementation 110% of the V_o size), and generates grain objects at random points inside the volume $V_b - V_o$. Grain positions are chosen by randomly choosing two vertices in the volume $V_b - V_o$ then linearly interpolating between them by a random amount. Grains are repositioned every 20 ms as recommended by Templin et al. [2]. Each grain is rendered in world space, meaning the user's eyes can focus on individual grains.

We implemented the two rendering techniques in four different styles. The first two use *Shader* rendering, applying the Perlin shader to either the user's view or objects in the scene. *Headset Grain* (Figure 2a) renders the grain texture on a plane covering the user's entire field of view, anchored to the user's head about 5 cm away. Users perceive film grain particles along a single "sheet" in their vision. *On-Surface Grain* (Figure 2b) uses the Perlin shader to render grain as a flat texture, which is then added to the surfaces of objects in the scene.

The last two styles use *Volumetric* rendering, creating grains as individual 3D particles in the scene. *Object Cloud* (Figure 2c), inspired by previous work on perceptually-motivated stereoscopic film grain [2], renders grains at varying density depending on the camera's distance from objects in the scene. Grains are randomly distributed in space where there is no object in view, but become more densely concentrated around objects in the scene depending on their depth from the camera. Infinite-depth grain is achieved by creating one large surrounding grain volume for particles to appear randomly, while objects in view dynamically add or remove particles from their grain volume V_b depending on their Euclidean distance to the headset. *Free Cloud* (Figure 2d) disperses grains randomly within a large predetermined cylindrical volume surrounding the user. Because the volumetric grain technique is applied to the user instead of an object, we set the size of V_o to be 0 such that grains can appear at any point within the user's surrounding cylindrical volume.

An 8-participant pilot study of these lens flare and film grain techniques showed that participants preferred volumetric film grain effects and distant lens flare effects in VR environments, but context and the surrounding scene can affect the strength of this preference.

[1] Matthias Hullin, Elmar Eisemann, Hans-Peter Seidel, and Sungkil Lee. Physically-based real-time lens flare rendering. *ACM Trans. Graph.*, 30(4), 2011.

[2] Krzysztof Templin, Piotr Didyk, Karol Myszkowski, and Hans-Peter Seidel. Perceptually-motivated stereoscopic film grain. *Comput. Graph. Forum*, 33(7):349–358, October 2014.