Character Rendering for VR in "Blood & Truth"

James Answer
Sony Interactive Entertainment London Studio
London, United Kingdom
james.answer@sony.com

Figure 1: Screenshots from Blood & Truth for PlayStation VR.

ABSTRACT

The performance challenges of real-time rendering for VR are well documented - this poses a particular challenge for rendering realistic human characters, with many of the usual techniques popularised in games not scaling well to high pixel counts and framerates. For Blood & Truth on PlayStation VR, we adapted existing techniques to fit our VR focused forward renderer, and invented novel replacements for expensive screen-space effects.

CCS CONCEPTS

• **Computing methodologies** → **Rendering**; **Rasterization**; **Virtual reality**.

KEYWORDS

games, VR, real-time, characters

ACM Reference Format:

James Answer. 2020. Character Rendering for VR in "Blood & Truth". In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH '20 Talks), August 17, 2020.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3388767.3407330

1 SKIN RENDERING

Many 2D games currently use some variation of screen-space subsurface scattering to reproduce the scattering profile of skin. This is undesirable for us, as screen space effects do not scale well to our resolution requirements. Instead, we have used the pre-integrated skin shading model described in [Penner and Borshukov 2011]. This allows us to acceptably reproduce the main characteristics of skin in a single pass in our forward renderer. In addition to the published technique, we also use a per-channel fitted curve rather than a lookup texture, which allows us to reduce the number of texture fetches required.

The skin lighting was further improved by the use of pre-calculated bent normals for base and wrinkle poses, which allow us to reduce light leaking from ambient lighting. We can also use the light direction, bent normal and ambient occlusion terms to construct an ad-hoc shadowing factor as in [Brinck et al. 2016] to add additional shadow detail which can only be captured through our textures rather than the necessarily low poly geometry and low resolution shadow maps we use.

The bent normals and wrinkle maps on the face required many texture reads in the face shader. To reduce the cost of this, only the base normal and albedo maps were stored at full resolution, with any wrinkle maps and bent normals being stored as an offset from the base pose at a lower resolution, as we only needed coarse changes for these maps, rather than the pore detail.

Our facial animation is entirely driven by joints, which requires the normals to be recalculated per-frame in a geometry shader, as the standard skinning methodology of just rotating the normals with the joint matrix does not accurately describe the deformed surface when translations are used. When calculating the new normals, getting the area of the deformed triangle is trivial, so we can compare it against the area of the rest pose and get some idea of how the skin is squashing and stretching. Inspired by the work in [Nagano et al. 2015], which measured the effects on microsurface roughness based on the deformation, we apply an ad-hoc artist driven offset to the skin roughness based on the area difference.

SIGGRAPH '20 Talks, August 17, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-7971-7/20/08.
https://doi.org/10.1145/3388767.3407330

2 JOINT BASED ANALYTIC EFFECTS

Features such as eyes and teeth have historically been difficult to accurately light in games. Methods such as screen space ambient occlusion (SSAO) and reflections, as well as increased shadow resolution have improved this over the last few years, but they are often employed for in very constrained conditions for non-interactive cut-scenes, and the performance requirements for some of these features would be prohibitive for us. Even where we do have them, like in the case of SSAO and shadows, the resolution is insufficient to capture the detail needed. We observed that many of the lighting features present on eyes and teeth can be approximated using existing facial skeleton joint positions to estimate them. By passing through these positions to the specialised eye and mouth shaders as shader constants, we can construct approximations to a distance field of the surrounding geometry, which can be used to modify a number of material parameters. These include ambient occlusion, bent normals, soft tissue transitions, tear lines and approximating an occluder plane to generate eyelash reflections.

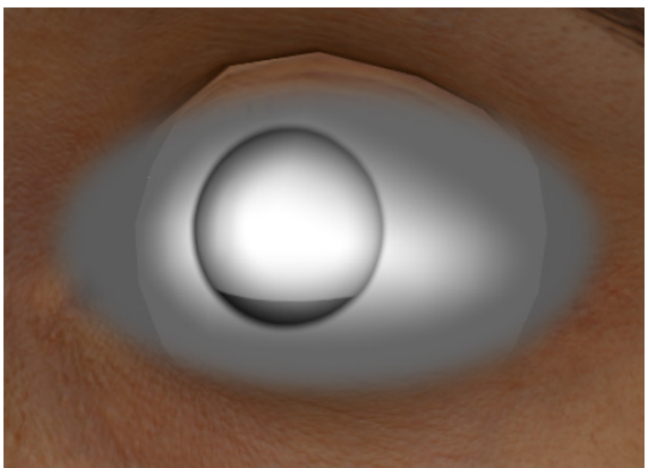

Figure 2: Ambient occlusion generated from joint positions.

This technique scales well to high resolutions, and has the additional advantage that it works even if only a small number pixels of the target geometry are visible, where screen space techniques may fail.

3 "STICKY" CATCH LIGHTS

A common technique used in photography and cinema is to give the eyes more life through the use of "catch lights", which add a small, bright specular highlight, placed within the iris where it is most visible. Game cinematics often employ the same techniques, but in VR this can be more difficult to achieve with moving characters and without camera cuts to hide lighting rig changes. We implemented a simple solution in the eye shader to bias existing scene lights towards the view vector when it gets to a physically plausible range. This gives a more visually pleasing highlight, without noticeable discontinuities and no additional setup time or specially linked lights, and more freedom in creating light and shadows on the face without worrying about maintaining the highlight.

4 HAIR RENDERING

To avoid sorting artifacts and reduce cost, we used a temporal stochastic dither on alpha tested cards, using temporal anti-aliasing to give the impression of transparency. This had the side effect of sometimes producing flickering artifacts in the user's peripheral vision due to the nature of the reduced resolution we render there, so to mitigate this we blend to a simple threshold in the outer extents of the view.

5 RESULTS

Blood & Truth achieves comparable quality of character rendering to many traditional 2D games, but using techniques that allow us to render at higher resolutions, and maintain constant 60 frames per second rendering.

ACKNOWLEDGMENTS

The Sony Interactive Entertainment London Studio technology and character teams for the engine and art, and 3Lateral for the facial capture.

REFERENCES

Waylon Brinck, Andrew Maximov, and Yibing Jiang. 2016. The Technical Art of Uncharted 4. In *ACM SIGGRAPH 2016 Talks (SIGGRAPH '16)*. Association for Computing Machinery, New York, NY, USA, Article 60, 1 pages. https://doi.org/10.1145/2897839.2936731

Koki Nagano, Graham Fyffe, Oleg Alexander, Jernej Barbiç, Hao Li, Abhijeet Ghosh, and Paul Debevec. 2015. Skin microstructure deformation with displacement map convolution. *ACM Transactions on Graphics* 34, 4 (Aug. 2015). https://doi.org/10.1145/2766894

Eric Penner and George Borshukov. 2011. *GPU Pro 2* (1st. ed.). A K Peters/CRC Press, New York, Chapter 2.1, 41–55. https://doi.org/10.1201/b11325