Time and Memory Efficient Displacement Map Extraction

Vinod Melapudi
vmelapudi@blizzard.com
Blizzard Entertainment
USA

Parag Havaldar
phavaldar@blizzard.com
Blizzard Entertainment
USA

ABSTRACT

Displacement maps are used to enhance the final rendered animation at most high-end studios including Blizzard Entertainment. However, commonly available products (Mudbox, Zbrush) that are used to sculpt high resolution models and extract displacements from low resolution counterparts are not conducive for quick iterative workflow that artists normally demand. Here we will summarize the drawbacks in accommodating such commercial products in our pipeline. We then present our custom solution developed using open industry standard libraries like OpenSubDiv (Pixar), OptiX (NVIDIA) and Embree (Intel). We also highlight how we employed frequency analysis (Discrete Cosine Transform) to extract time efficient and memory optimal displacement samples, all of which has objectively improved overall productivity in our workflows involving displacement maps.

CCS CONCEPTS

• **Computing methodologies → Rendering**.

ACM Reference Format:

Vinod Melapudi and Parag Havaldar. 2021. Time and Memory Efficient Displacement Map Extraction. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH '21 Talks), August 09-13, 2021.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3450623.3464672

1 INTRODUCTION

In a production pipeline, modelers normally build lower resolution assets (Figure 1) for layout, rigging, and animation steps, that allow artists to interact with the asset at manageable frame rates in commonly availed DCCs (e.g. Maya). The modelers additionally sculpt details on a much higher resolution mesh (Figure 1). These sculpted details are extracted and stored as high-resolution displacement texture maps. During final rendering, displacement shaders use these displacement maps to displace the view based subdivided mesh, thus restoring the sculpted geometric details while rendering. These displacements can be extracted either as a 3-channel vector-displacement or a single channel normal-displacement. Here we focus on normal-displacement maps which encode mapped distances from the low-resolution surface to the sculpted high-resolution surface.

At our studio, we originally engaged off-the-shelf tools like Mudbox (Autodesk) and ZBrush (Pixologic) to extract displacements. This posed many challenges to our pipeline. To extract maps, artists had to iterate numerous times varying parameters that controlled the quality of the extraction, resulting in displacement maps with high memory footprints. Artist workflows and directorial feedback often changed a few shapes/parts of the complex model, but the extraction process mandated the extraction for all the shapes on each iteration. Each iteration was time consuming. Also, extractions were only possible in interactive mode and could not be procedurally run on a farm. Furthermore, such off-the-shelf software provided limited API interfaces that prevented us from customizing or optimizing the process to be more time and memory efficient. These limitations led us to develop our own custom extraction tools addressing the same.

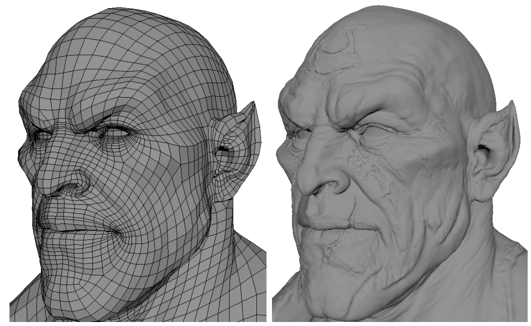

Figure 1: Low resolution base geometry and High resolution sculpted geometry

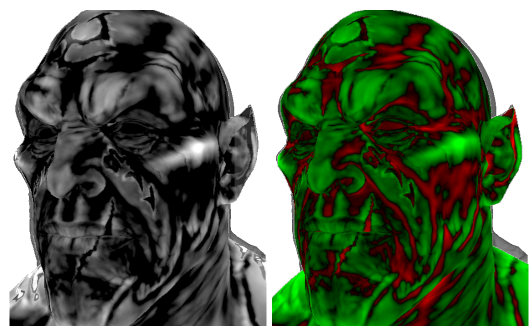

Figure 2: Displacement map in grey scale and with Positive(green) / negative (red) displacements

2 APPROACH

Our extraction tool takes a production asset's base mesh and a higher resolution sculpted mesh as inputs and produces a Ptex texture with scalar displacement values. Employing the Ptex architecture, we process one base mesh face at a time to extract the displacement data. For each face we need to determine the optimal number of samples (texels) required to reflect the fidelity of the corresponding sculpted details. We establish this optimal resolution by a two-step process. First, we extract displacements at predetermined over-sampled uv locations. Second, we employ frequency analysis over the displacement values to determine the minimal

Table 1: Extraction time and Ptex size comparison

Asset	Num. polygons (in 1000s)		Evaluation time (seconds)			Total Texels (millions)	
	Base	Displaced	Mudbox	OptiX	Embree	Non-optimized	DCT optimized
Cape	4K	4215K	69.0	8.6	7.5	17	7
Head	8K	7970K	82.0	17.5	15.8	32	11
Body	17K	16049K	149.0	28.9	26.7	69	30

Note: Extraction time in Mudbox was approximately determined from UI interaction.

Figure 3: Render without and with displacements

resolution. The predetermined resolution is chosen, based on the subdivision level of the sculpted geometry. During extraction, the precomputed texel resolution is used to generate uniformly sampled uv locations across each base mesh face. We use the OpenSubDiv library to subdivide the base mesh and evaluate the limit surface positions and normals for each of these uv samples. Displacements are extracted from these limit surface locations by intersecting the rays along normal direction with the high-resolution sculpted geometry. To build the acceleration structure for high-resolution geometry and process ray tracing, we adapt NVIDIA's OptiX and Intel's Embree which evaluate on GPU and CPU, respectively. This allows our artists to make a computation choice depending on availability of GPU resources when working interactively or when batch processing on the farm.

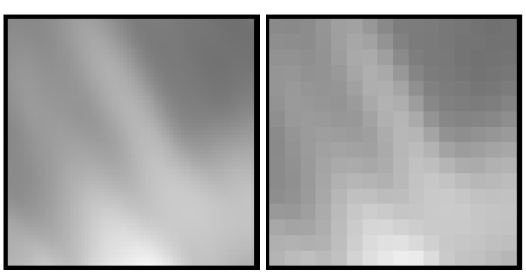

Figure 4: Non-optimized 64x64 and optimized 16x16 texels displacements of a face

Additionally, we need to ensure that the displacement maps are memory optimal, to allow for minimal storage requirements and favorable load-save times. So, extraction process is followed by frequency space analysis on the over-sampled displacement data (Figure 4). The over-sampled data is converted to Fourier space representations to identify the highest nonzero frequency coefficients present. If f(x,y) is used to define the over sampled locations at which displacements are computed then F(u,v) give the frequency coefficients for the extracted samples as follows [1]

$$F(u,v) = \left(\frac{1}{4}C(u)C(v)\right)\left[\sum_{x=0}^{x=7}\sum_{y=0}^{y=7} f(x,y) \times \cos\frac{(2x+1)u\pi}{16}\cos\frac{(2y+1)v\pi}{16}\right]$$

$where: C(u), C(v) = 1/\sqrt{2}$ for u,v = 0 and C(u), C(v) = 1 otherwise

Given that, the classical sampling theory suggests, atleast twice the maximal frequencies in u and v are needed to capture all the details in the spatial domain, we find the maximal u and maximal v in our F(u,v), for which F(u,v) is not zero. Twice the u (and v) value then suggests, the number of samples required in the x and y direction respectively, noting that the number of samples that span in the x and y directions may not be the same. Given these resolution limits, we re-sample each face's displacement data appropriately to optimal lower resolutions (Figure 4). This approach has shown to reduce the memory footprint to approximately 60 percent in our production assets, as reflected in the total texels in Table 1. It also improved upon previous workflow, where artists had to iteratively extract in interactive mode to obtain an optimal map.

The displacement map capturing the sculpted details are shown in grey scale and green-red scale reflecting positive-negative displacements in Figure 2. Figure 3 shows the base mesh rendered without and with displacements. We have tested the maps with both Renderman and Redshift renders, which are primarily used at our studio. Our approach leveraging OpenSubDiv, OptiX and Embree performs better than the solutions available through Mudbox and ZBrush and we compare the evaluation times in Table 1.

3 CONCLUSION AND FUTURE WORK

We have leveraged industry standard libraries, which are performance tuned for ray tracing on GPU and CPU, to extract displacement maps reflecting artist sculpted details. We have enabled the extraction process to be run on our farm, which artists leverage to batch process multiple assets. Our approach has proved to be both time and memory optimized compared to off-the-shelf tools. Since 2019 our productions have benefited by this technology in accelerating the workflow. For future work, we have recently developed an alternate extraction process leveraging subdivision mapping between base and sculpted high-resolution mesh. This allows to extract vector displacements which is yet to be pipelined and demonstrated in production.

[1]The JPEG Still Picture Compression Standard –Gregory K Wallace, Communications of the ACM, April 1991