Temporal and Spatial Anti-Aliasing for Rendering Reflection on a Water Surface

Namo Podee
Hokkaido University

Nelson Max
University of California, Davis

Kei Iwasaki
Wakayama University

Yoshinori Dobashi
Hokkaido University

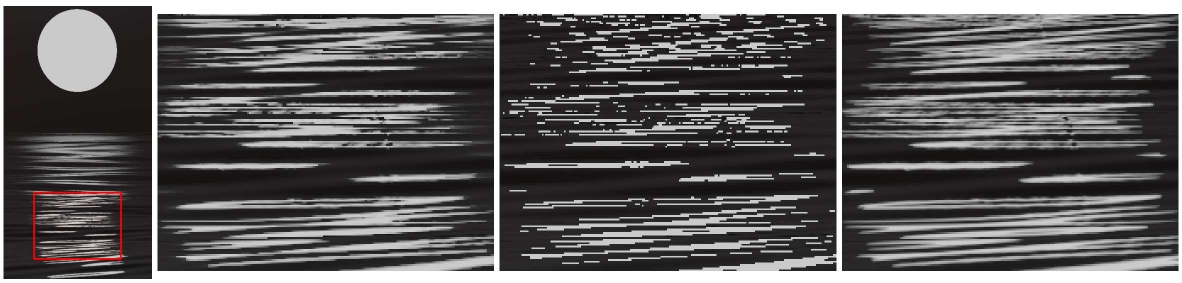

(a) (b) (c) (d)

Figure 1: (a) The result of our method, (b) the closeup views of our method, (c) without anti-aliasing image, (d) reference image. The environment map "Milkyway" by Blochi, via sIBL Archive (www.hdrlabs.com/sibl/archive.html)

CCS CONCEPTS

• **Computing methodologies → Rasterization**; Reflectance modeling.

KEYWORDS

computer graphics, real-time rendering, anti-aliasing, reflection

ACM Reference Format:
Namo Podee, Nelson Max, Kei Iwasaki, and Yoshinori Dobashi. 2019. Temporal and Spatial Anti-Aliasing for Rendering Reflection on a Water Surface. In *Proceedings of SIGGRAPH '19 Posters*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3306214.3338599

1 INTRODUCTION

The ocean surface is highly dynamic. It moves rapidly and thus its shading changes rapidly as well. Usually, this doesn't pose any problems if the shading is smooth. However, for a surface that has a strong highlight or bright reflection moving rapidly, it causes an inaccurate and unnatural flickering. In the traditional rendering algorithms, each frame is rendered independently at a discrete time, resulting in serious temporal aliasing artifacts. Particularly, for a wavy water surface, reflection vectors may not hit the light source even though they actually hit for part of the frame time. Removing such aliasing in real-time is an active research area and many methods have been proposed [Jimenez et al. 2011]. They can improve the fidelity and efficiency of the rendering method. However, their focus is on spatial anti-aliasing and most of them do not address the temporal aliasing problem, particularly the one observed in rendering a reflected image of a light source on the water surface.

In this paper, we present a method that can remove the spatial and temporal aliasing simultaneously. The basic idea is to compute the intersection of the light source with a plane formed by two reflection vectors for neighboring frames. This provides us with a fraction of time when the light source is visible on the water surface. We combine this idea with a traditional spatial anti-aliasing method.

2 OUR APPROACH

Our goal is to compute the contribution of a spherical light source over a period of time between rendering frames. We assume that the water wave is represented as a sum of sine waves with different frequencies and directions. For each pixel, our method first decomposes the water wave into two frequency bands: low and high-frequency bands, namely non-aliasing waves and aliasing waves, respectively. The high-frequency components, or aliasing waves, cause spatial aliasing due to under-sampling. We develop two methods, non-aliasing wave rendering and aliasing wave rendering, for these two frequency bands, respectively. The process for both methods is the same but the spherical light source is blurred to account for the effects of the BRDF caused by the aliasing waves.

2.1 Aliasing detection

We use the clamping anti-aliasing method [Norton et al. 1982] to decompose the water waves into aliasing and non-aliasing waves.

SIGGRAPH '19 Posters, July 28 - August 01, 2019, Los Angeles, CA, USA

ACM ISBN 978-1-4503-6314-3/19/07.
https://doi.org/10.1145/3306214.3338599

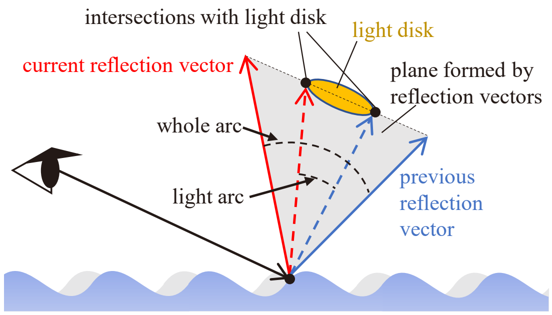

Figure 2: Basic idea of our method.

According to sampling theory, to avoid aliasing the sampling frequency must be higher than the twice the highest frequency of the water wave. In our case, we calculate the projected wavelength of each sine wave on each screen pixel. Each of the sine waves is then classified into either the non-aliasing and or the aliasing waves according to the Nyquist frequency, which is $\frac{1}{2\sqrt{2}}$ pixels. Any projected wave that has lower frequency is a non-aliasing wave and vice versa for aliasing wave. We use soft classification, with a smooth amplitude transition starting before the Nyquist limit, to avoid arcs in the image of sudden appearance changes.

2.2 Non-aliasing wave rendering

For the non-aliasing waves, we sample a single point on the water surface corresponding to the pixel center and compute the contribution of reflected light over the time interval between the current and the previous frames. We assume that the light source is far distant from the sample point and is approximated by a disk facing toward the sample point. Two reflection vectors are computed by using the normal vectors at the previous and the current frames. We then calculate intersection points between the light disk and a plane formed by the two reflection vectors, as shown in Fig. 2. A reflected viewing ray at the sample point hits the light disk when it lies between the directions to the two intersection points. Thus, the fractional contribution of the light source between the frames is obtained by the ratio of the angle of the light arc to the angle of the whole arc (see Fig. 2).

2.3 Aliasing wave rendering

For aliasing waves, a single sample point per pixel is not sufficient. We have to compute the average intensity of the reflected light over the pixel area taking into account the distribution of the surface normals. A straightforward solution is to generate multiple rays for each pixel, which significantly increases the computation time. Instead, we borrow the idea of the LEAN mapping technique [Olano and Baker 2010] for efficient computation.

We first calculate the covariance matrix of the normal distribution function (NDF). The covariance matrix for the sum of the sine waves is obtained by accumulating the covariance matrix for each of the waves, which can be calculated analytically by integrating the normal direction of the whole sine wave.

Then we transform the NDF into the reflection space to obtain the reflection distribution function (RDF). The RDF represents the distribution of the reflection directions. We are inspired by [Zwicker et al. 2001], which transforms a Gaussian kernel from camera space to ray space by using a local affine approximation. We represent both NDF and RDF as elliptical Gaussians and use the local affine transformation to approximate the RDF. The convolution of the RDF and a light disk is the light contribution of each normal direction. To avoid temporal aliasing, we need to integrate a light contribution of a moving normal direction over time, and we assume that the normal direction to the non-aliasing waves is moving linearly over time. Then the temporal contribution of the convolution is computed by a line integral along the moving direction of the normal vector. However, it is costly to compute this in real-time.

To be efficient, we precompute the line integration over the convolution of the RDF and a light disk and save it into a texture with just three parameters: the perpendicular distance of the extended line segment to the center of the light disc, the distance along the line from the foot of this perpendicular to an endpoint of the segment, and the variance of the RDF. We approximate the elliptical Gaussian RDF as a circular Gaussian to reduce its dimension, which now needs only one variance parameter. Then the configuration becomes circularly symmetric, so only two parameters are needed for the line integral over the blurred disc.

3 RESULTS AND CONCLUSION

Fig. 1 shows that our method can significantly reduce aliasing artifacts. Please see the accompanying video for the animated version of this example. The reference image is created by generating an image with 64 times higher resolution and then downsampling it. For the temporal anti-aliasing, we generate 8 images between the neighboring frames and compute their average. The rendering time for our method and the reference images are 27 and 17214 ms respectively. These are measured on a laptop with Intel Core i7 @ 2.50Ghz, Memory 16 GB, and NVIDIA GeForce GTX 860M.

Our method reduces aliasing and increases the fluidity of wave reflection animation in real-time by using temporal and spatial anti-aliasing methods. It also deals with the changing position of a light source and works for any height/normal field, if its normal distribution is known for its aliased wave. However, by approximating the RDF as a circular Gaussian, our method loses accuracy for distant waves, which have more directionality. We are planning to address this issue by approximating the elliptical Gaussian with a set of circular Gaussians.

ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI Grant Number JP15H05924.

REFERENCES

Jorge Jimenez, Diego Gutierrez, Jason Yang, Alexander Reshetov, Pete Demoreuille, Tobias Berghoff, Cedric Perthuis, Henry Yu, Morgan McGuire, Timothy Lottes, Hugh Malan, Emil Persson, Dmitry Andreev, and Tiago Sousa. 2011. Filtering Approaches for Real-time Anti-aliasing. In *ACM SIGGRAPH 2011 Courses (SIGGRAPH '11)*. ACM, New York, NY, USA, Article 6, 329 pages. https://doi.org/10.1145/2037636.2037642

Alan Norton, Alyn P. Rockwood, and Philip T. Skolmoski. 1982. Clamping: A Method of Antialiasing Textured Surfaces by Bandwidth Limiting in Object Space. In *Proceedings of ACM SIGGRAPH 1982*. 1–8.

Marc Olano and Dan Baker. 2010. LEAN Mapping. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. 181–188.

Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. 2001. EWA Volume Splatting. In *Proceedings of the Conference on Visualization '01*. 29–36.