Adaptive Deferred Shading

Ian Mallett, Cem Yuksel, Amit Prakash, *University of Utah*

THE UNIVERSITY OF UTAH

i3D 2016

Reference

Adaptive Deferred Shading

Shaded Pixels

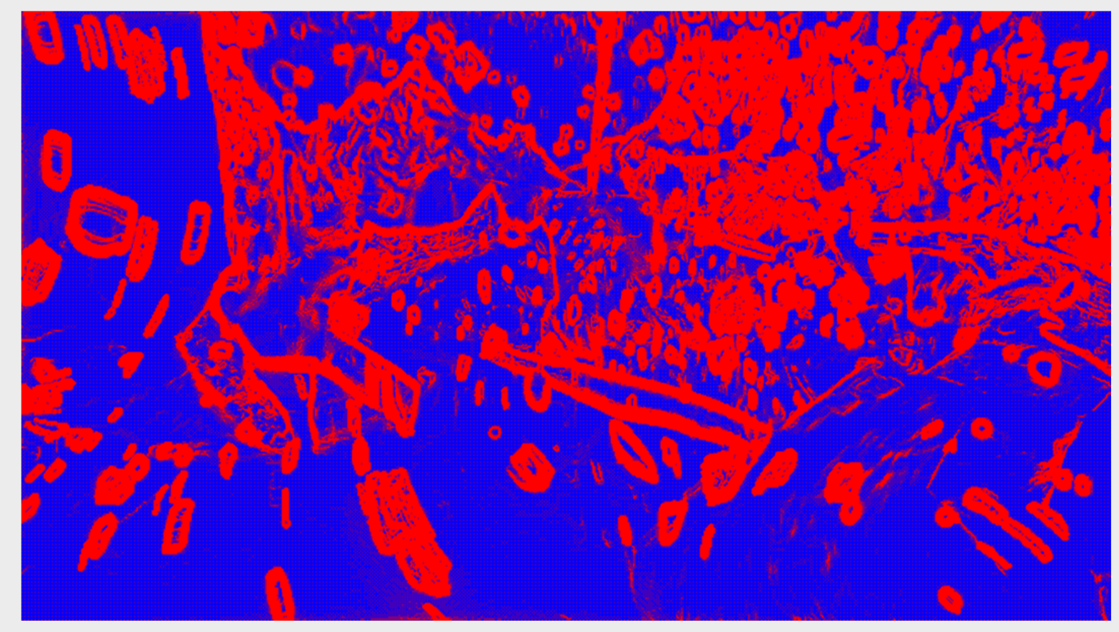

Scene from Unreal Engine "Elemental" demo reproduced using our technique. In this example, we achieve high quality (RMS 0.52, SSIM 0.999972), but require only 44.2% the original's shading rate. A visualization of the shaded pixels (red) shows how computation adapts locally to scene variance.

Abstract

The primary advantage of deferred shading is eliminating wasted shading operations due to overdrawing with forward shading. We propose a refactoring of the shading pass that allows further reduction in shading rate. Our adaptive shading algorithm decides whether to shade a pixel or to estimate the pixel color based on previously shaded pixels around it. Thus, we can produce the final image by shading only a fraction of its pixels, based on a threshold parameter. This provides a mechanism to tune performance and energy usage, both increasingly important issues as advanced graphics increasingly deploys to mobile devices. We simulate our method, which is geared toward motivating novel hardware, thoroughly on a wide range of real-world, game-quality scenes. We show that using our algorithm, even scenes that include high visual complexity can be generated by shading only a fraction of the pixels, with visually indistinguishable results.

Algorithm Overview

We begin with a coarse shading resolution and for each subdivision level we refine the shading rate progressively and recursively. At each $2^m x 2^m$ block of pixels, we use a threshold parameter to determine whether to invoke the shader or estimate the color of the pixel using the surrounding four.

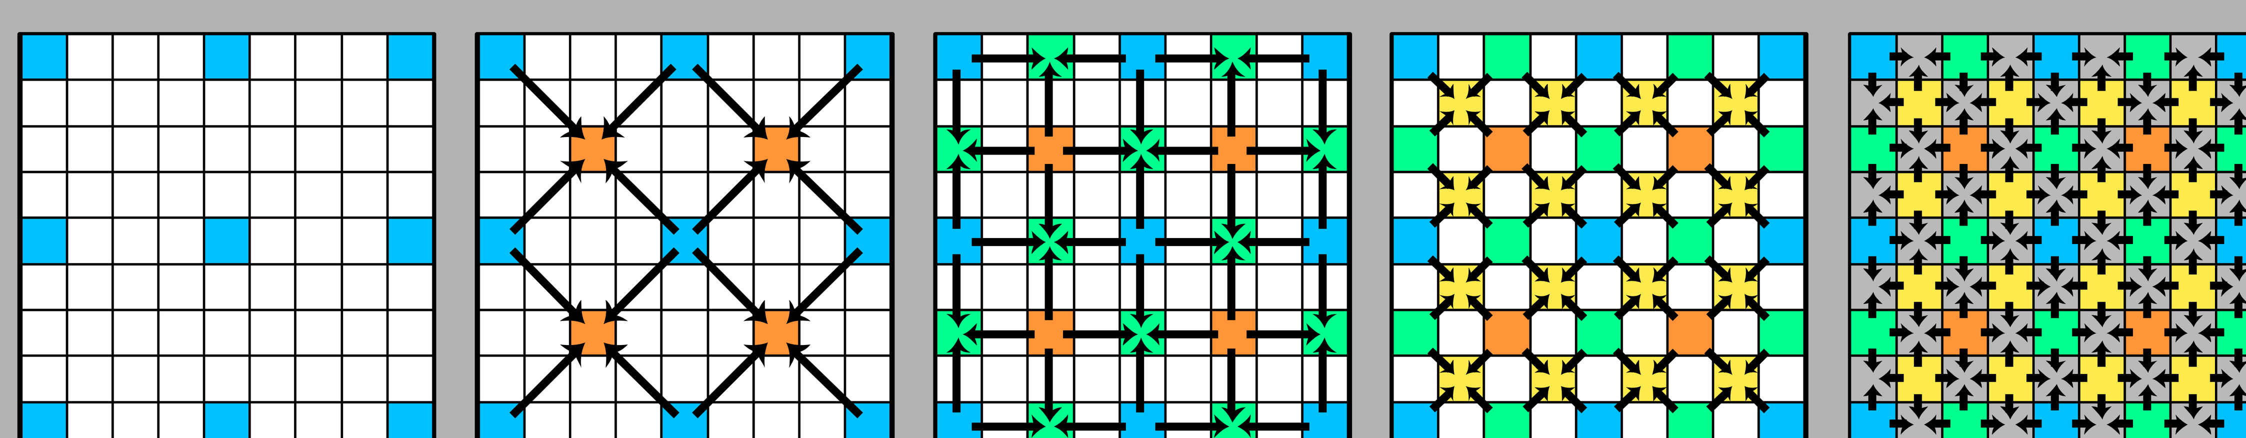

Test Results

We tested our algorithm on a variety of scenes from three ***Unreal Engine*** demos, which represent practical, real-world use-cases. Even under their challenging conditions, our method is both spatially and temporally stable. Typically, errors of our scheme are completely imperceptible, even in side-by-side comparison. When viewed in real-time, even lower shading rates can be achieved before any error can be detected.

Ours Low Threshold | **Ours High Threshold**

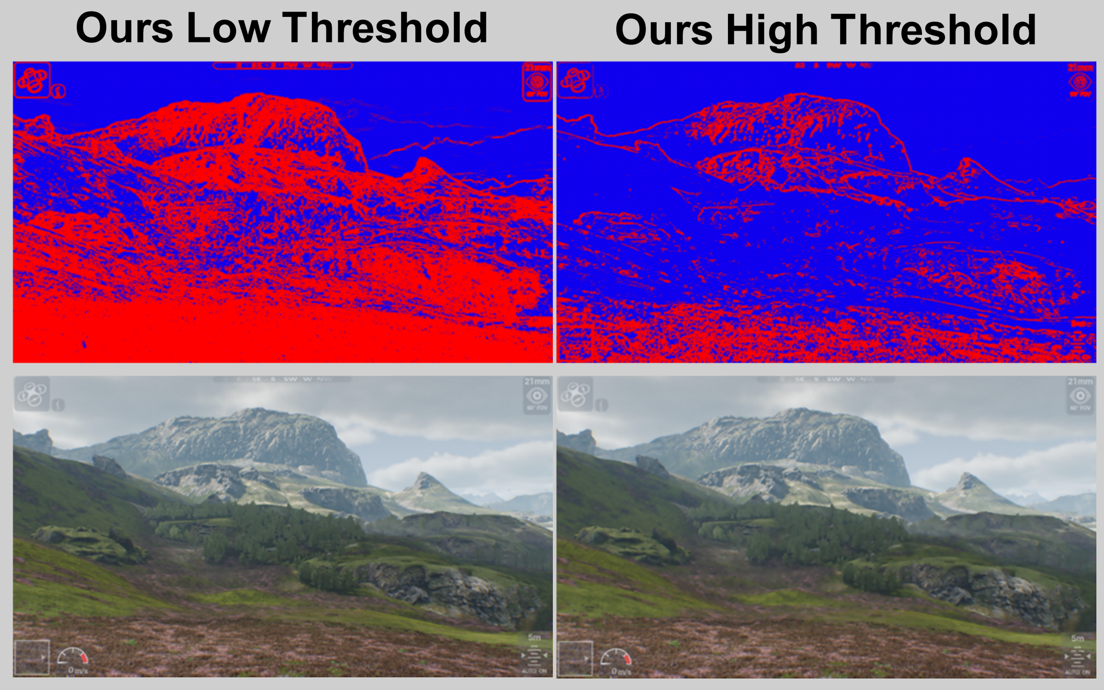

Reference | **Ours Low Threshold** | **Ours High Threshold**

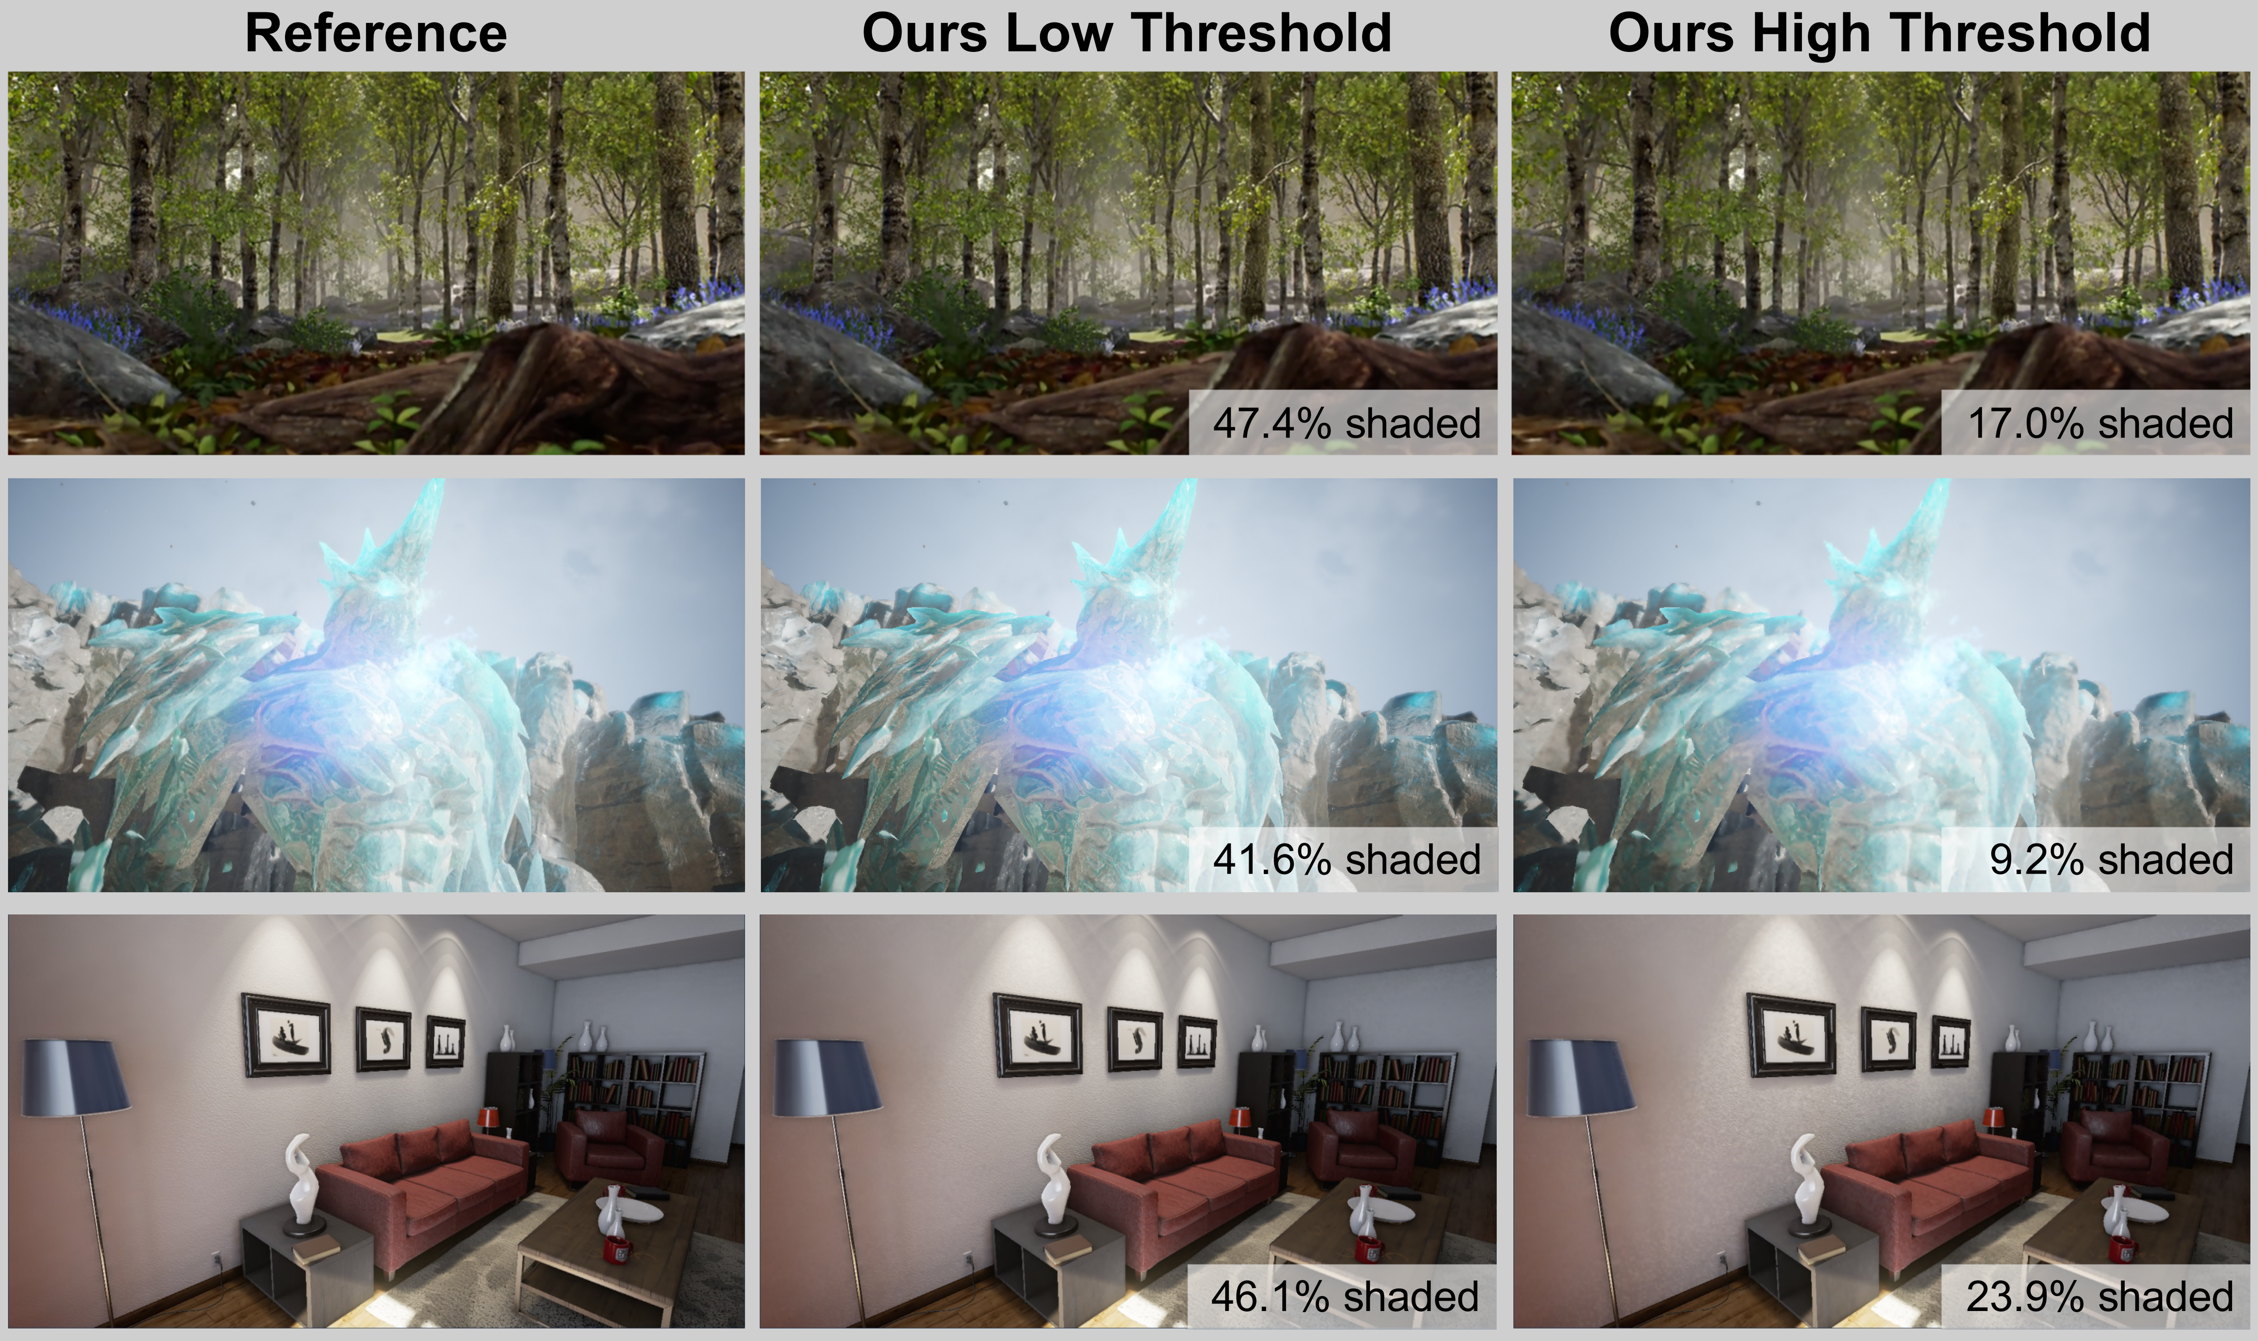

Deferred Shading Hardware

Our main goal is to motivate a hardware implementation for adaptive deferred shading. We propose an adaptive deferred shading unit that can act as a scheduler for fragment shading. This scheduler determines whether a pixel should be shaded or estimated using previously computed pixels. The threshold parameter can be controlled on-the-fly by the application or the system, as desired.

Bilinear Upscaling

Upscaling using bilinear interpolation produces unacceptably poor results.

Reference | **Ours** | **2x Upscale** | **4x Upscale**

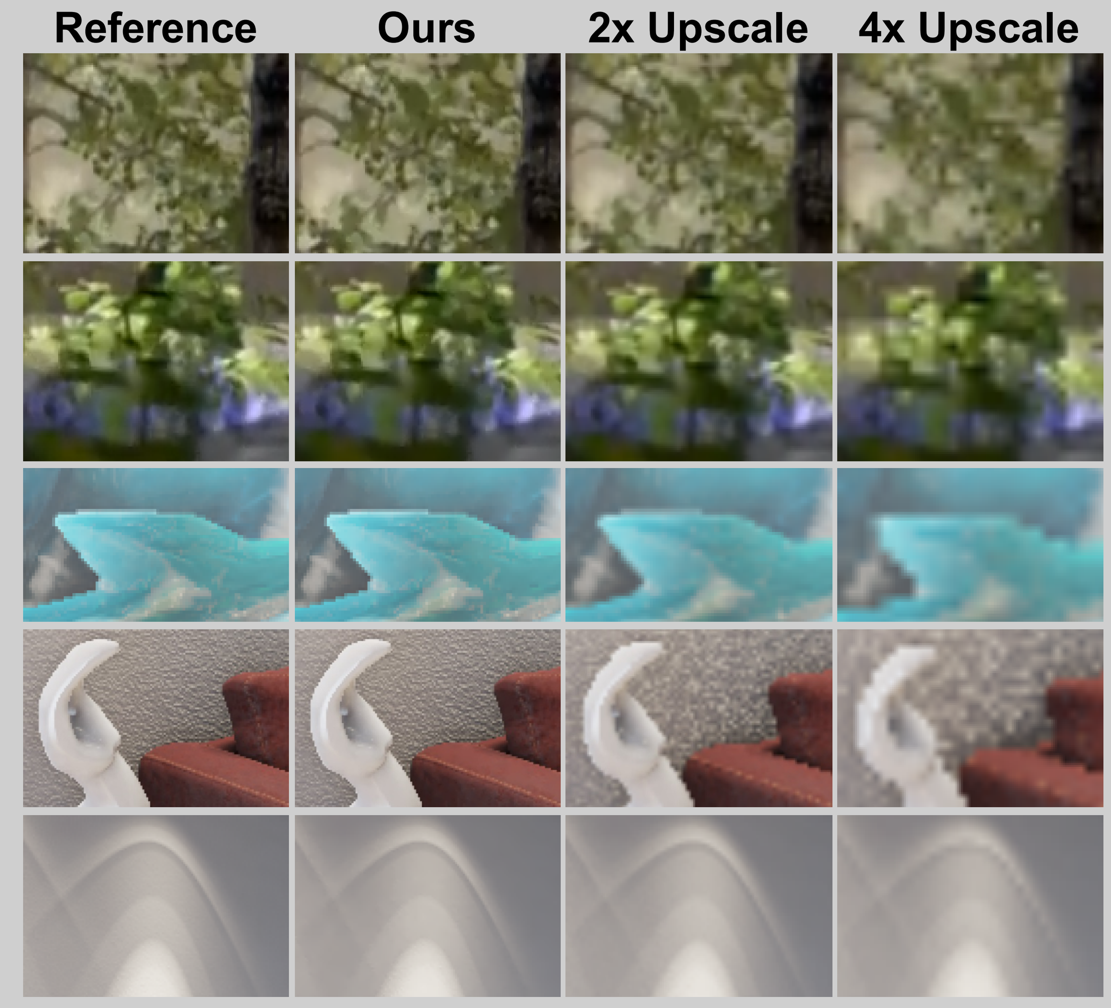

This work is supported in part by NSF award #1409129.