Approximating Dynamic G.I. in Image Space

-Tobias, Thorsten, Hans-Peter

Rohit Garg (05D26003)

Sriram Kashyap (08305028)

Motivation

- Realistic Shadows

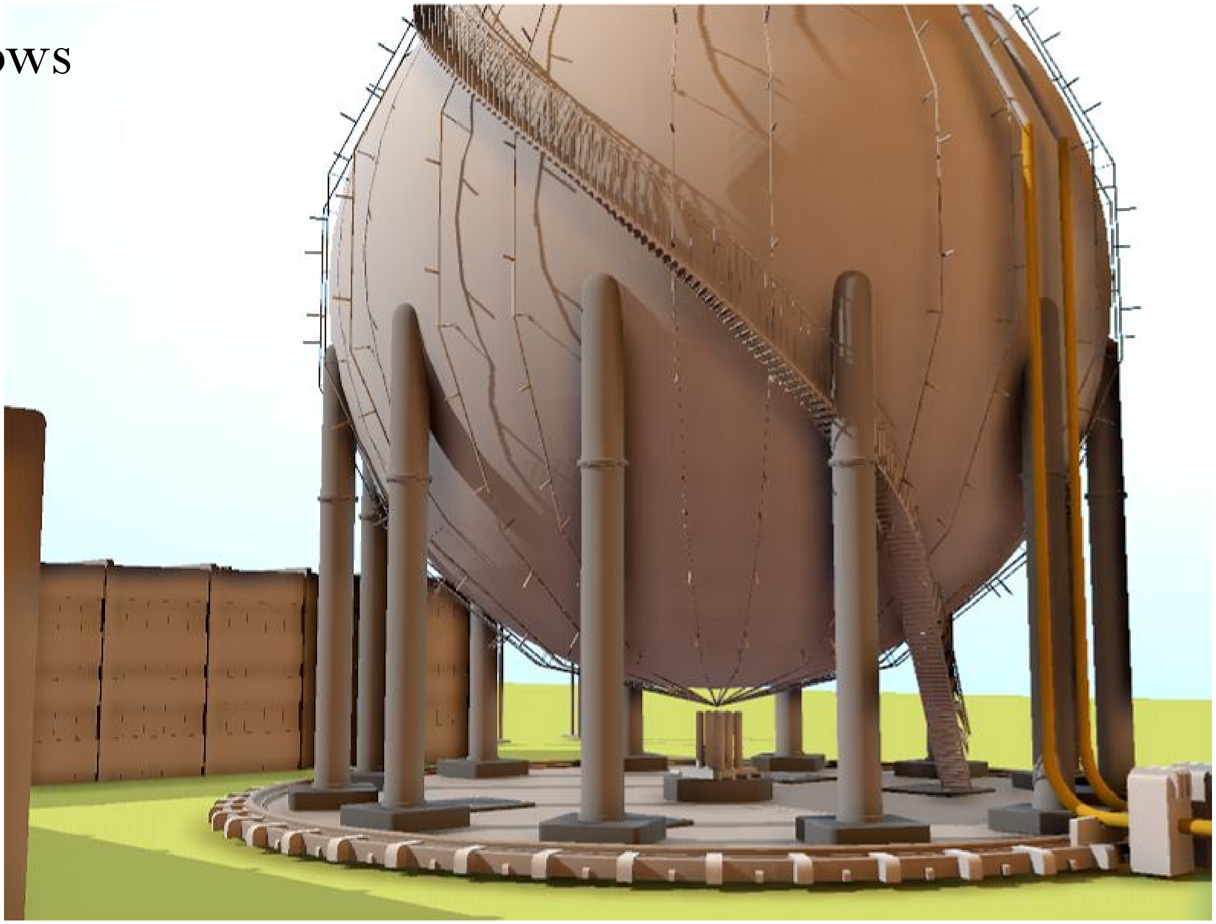

Motivation

- Realistic Shadows
- Color Bleeding

Ambient Occlusion

- Attenuation of ambient light due to local occlusion

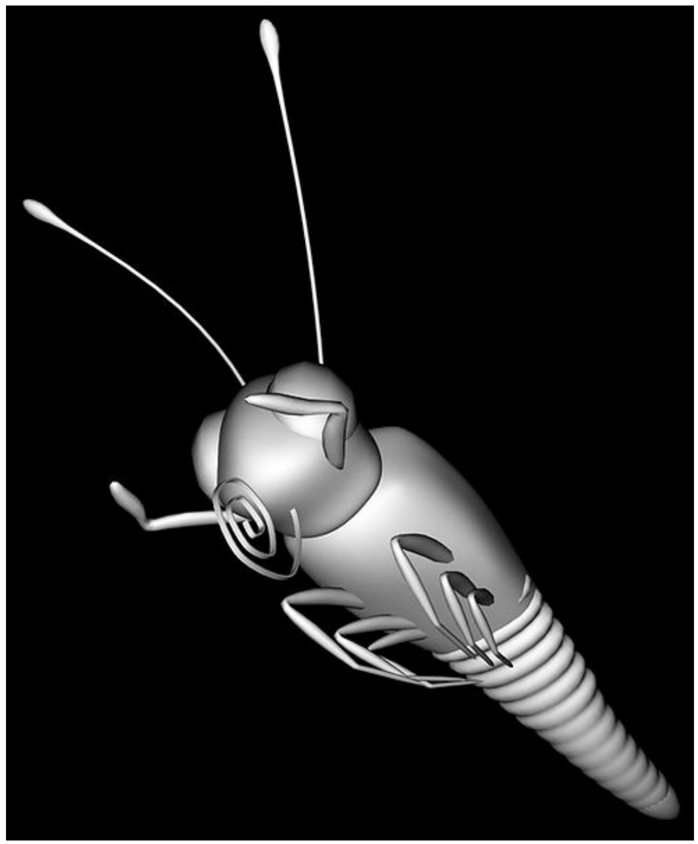

Diffuse Lighting

Ambient Occlusion

- Attenuation of ambient light due to local occlusion

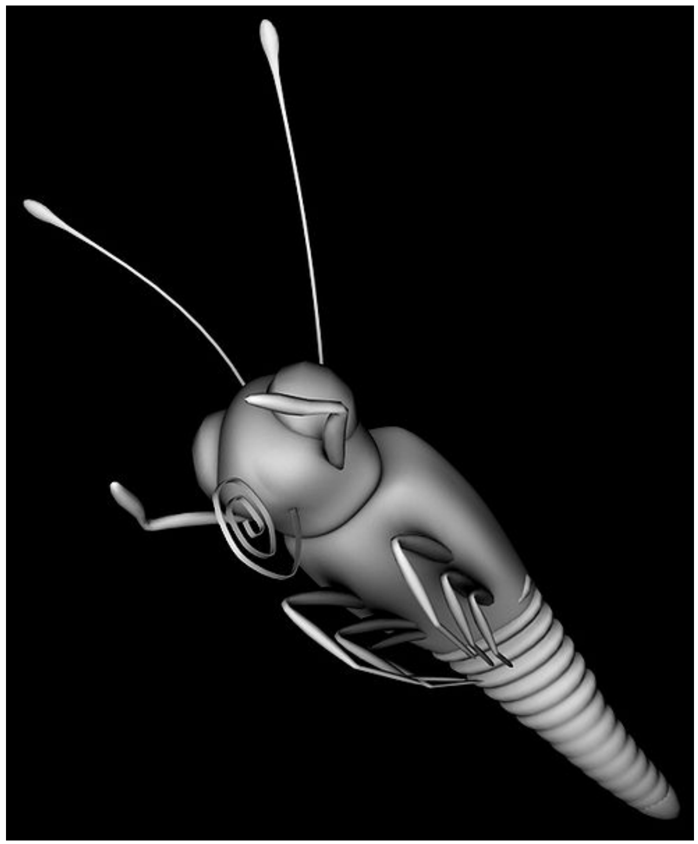

Ambient Occlusion + Diffuse

Basic Idea behind Occlusion

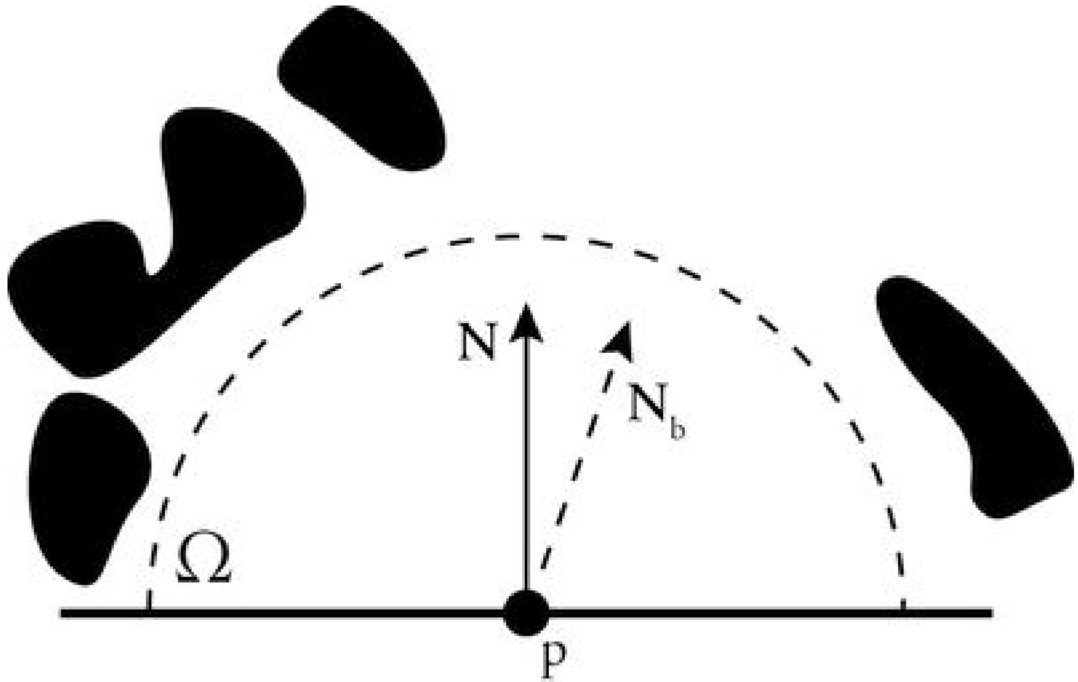

The occlusion at a point '*p*' on a surface with normal '*N*' can be computed by integrating the visibility function over the hemisphere '**Ω**' with respect to projected solid angle.

$$A_p = \frac{1}{\pi}\int_{\Omega} V_{p,\omega}(N \cdot \omega)\, d\omega$$

Ambient v/s Directional Occlusion

Ambient: Computes an average visibility value from all directions in the neighborhood. This occlusion value is then multiplied with the un-occluded illumination from all incoming directions.

Directional: Computes average incoming illumination. This accumulates incoming radiance from all un-occluded directions.

$$L_{\mathrm{dir}}(\mathbf{P}) = \sum_{i=1}^{N} \frac{\rho}{\pi} L_{\mathrm{in}}(\omega_i) V(\omega_i) \cos\theta_i \Delta\omega.$$

Screen Space Solution

Ambient Occlusion approximated in screen space (using a z-buffer)

Z-Buffer as a surface

Rasterization produces:

- 2D image (Frame Buffer)
- Depth map (z-buffer)

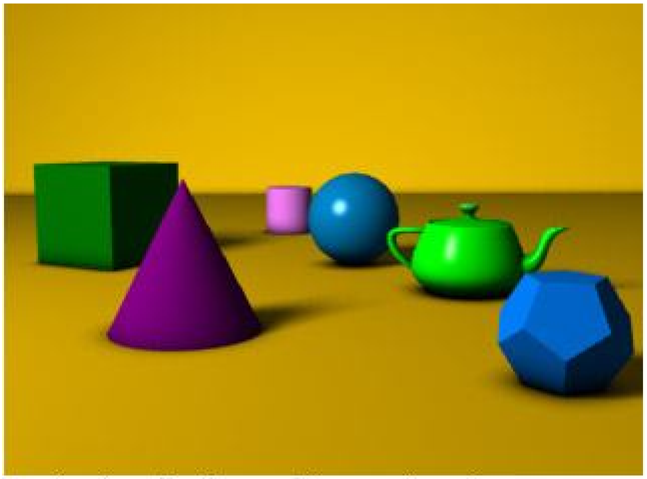

A simple three dimensional scene

The Z-Buffer can be thought of as a surface. It is a partial representation of the actual geometry of the scene.

Z-buffer representation

Z-Buffer as a surface

This means that any operations we do using the final rendered z-buffer are independent of the geometric complexity of the scene.

We highlight the utility of solving the occlusion problem in Image Space, as opposed to Object Space.

Here, we process only those pixels that are definitely visible.

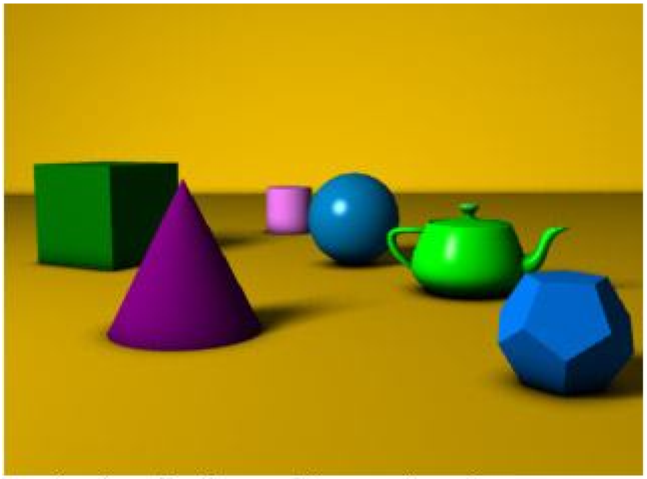

A simple three dimensional scene

Z-buffer representation

Directional Occlusion

- Randomly sample N points in the neighborhood of each pixel.
- The neighborhood is defined to be a hemisphere centered at that pixel, oriented in the direction of the local normal.
- Classify each sample point as ‘occluder’ or ‘visible’.
- Occluders are those points which lie below the z-buffer surface. Other points are visible.

Directional Occlusion

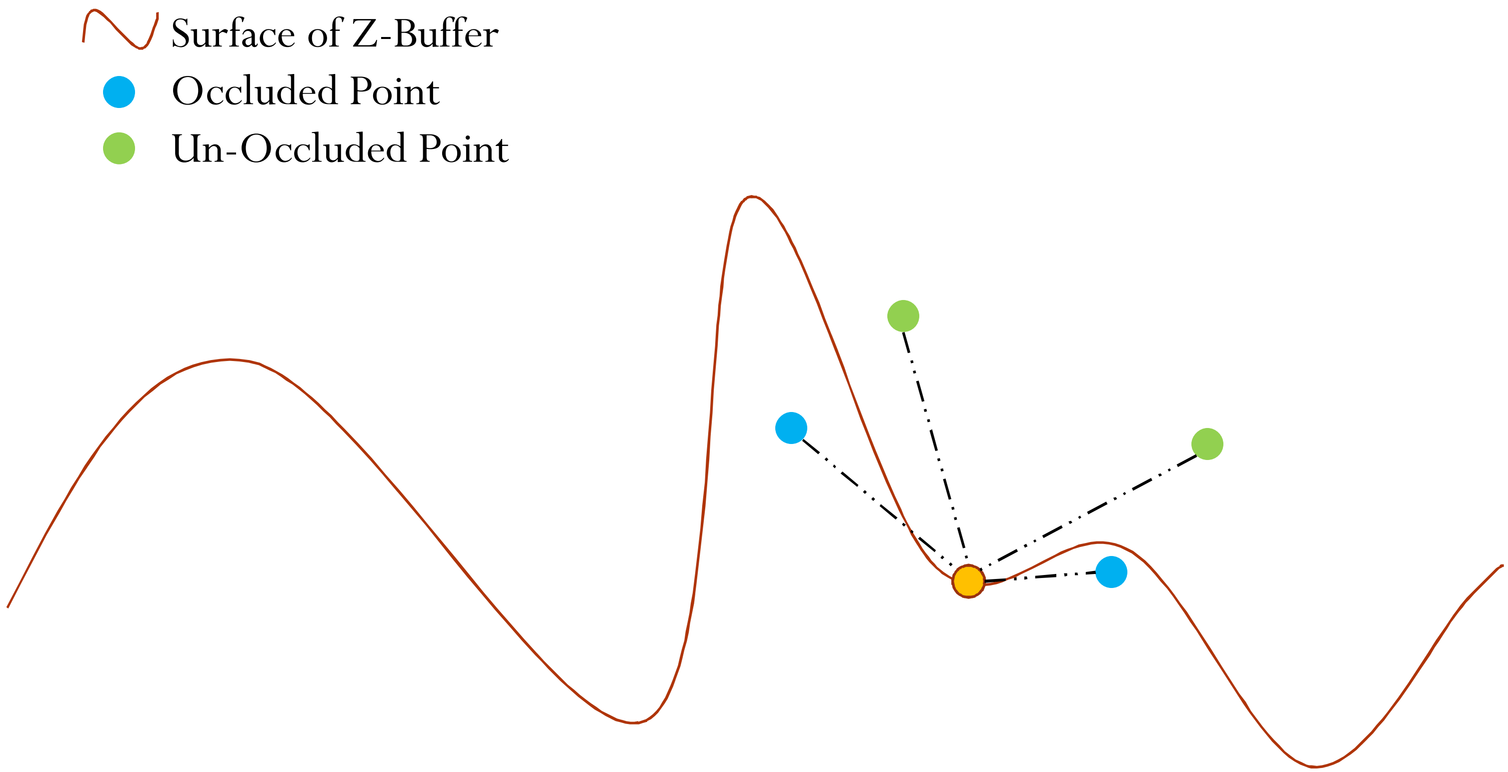

Directional Occlusion

- If an occluder exists along a direction, it means that no environmental light contribution can be received from that direction.
- The only light environmental contribution received at the pixel is from the directions of the visible points.

Directional Occlusion

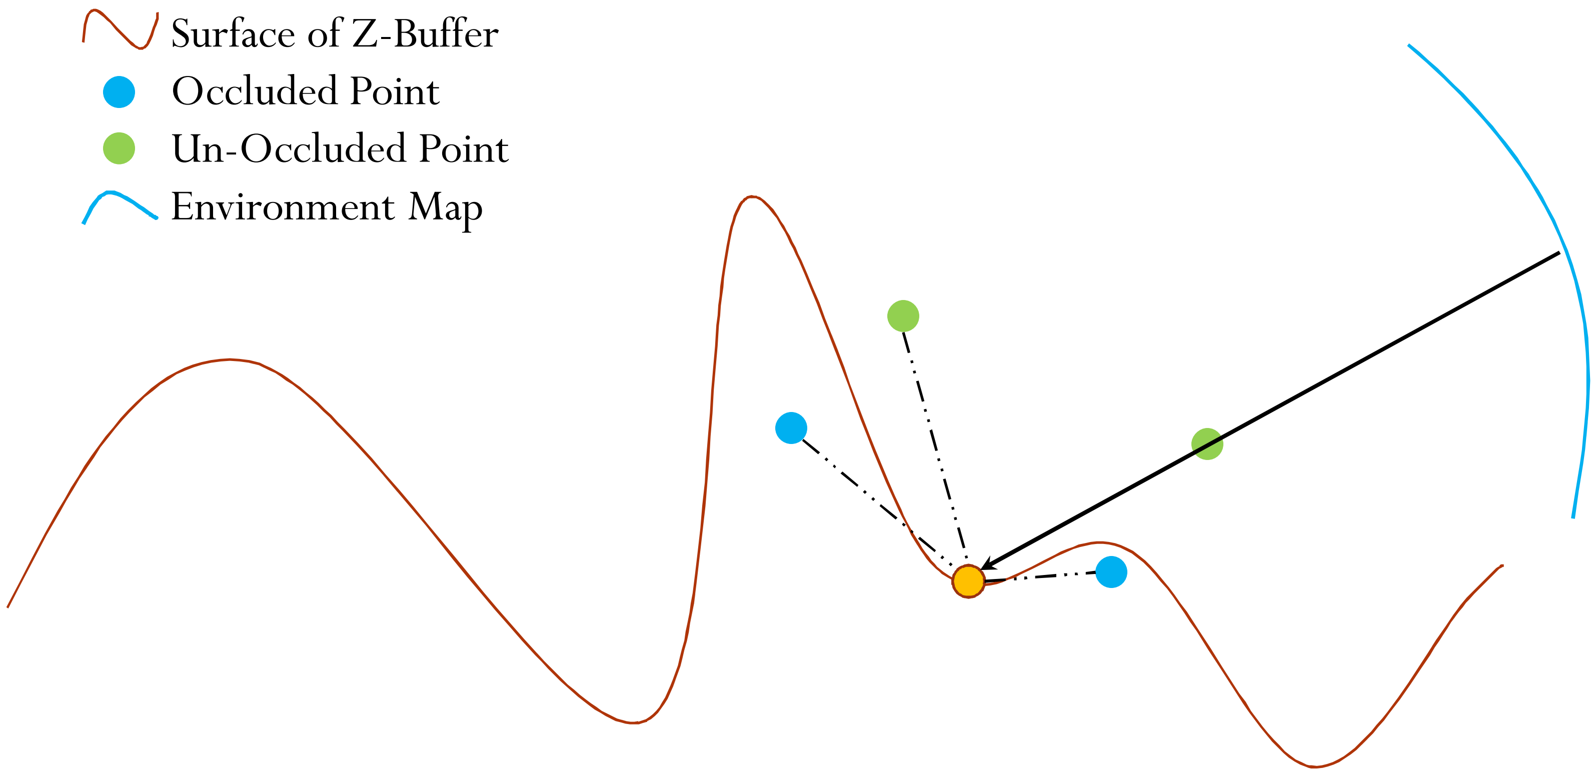

Environment Map

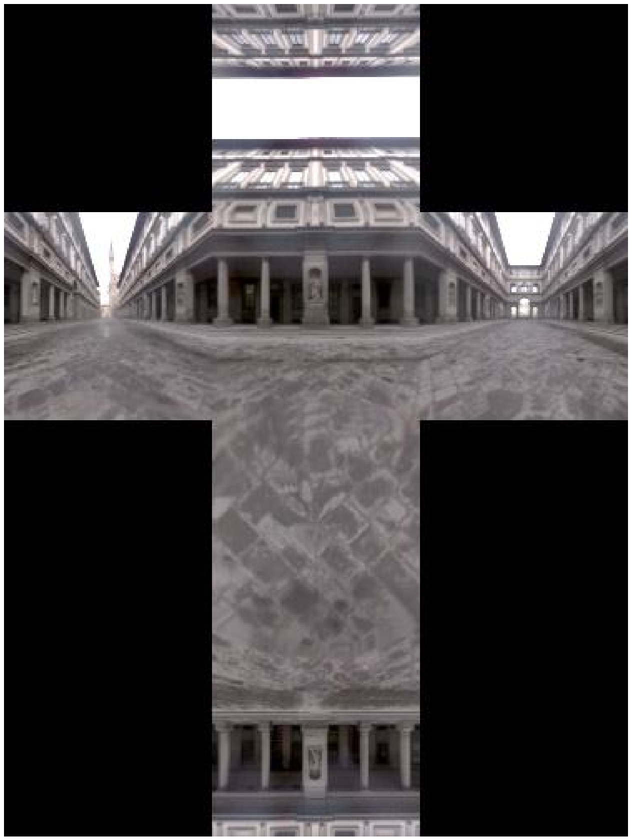

An environment map is a way of storing the surrounding environment. It is the information about color/light as seen at infinity from the current scene.

It completely encloses the scene. For example, an environment map can be constructed as shown here, by enclosing the scene in a large cube, and specifying the ‘texture’ as seen on each face of the cube.

Indirect Bounce

- The point on the z-buffer surface, corresponding to an occluder, can reflect some light towards the pixel.
- This computation is done in a manner similar to the form factor calculation of radiosity.

$$L_{\text{ind}}(\mathbf{P}) = \sum_{i=1}^{N} \frac{\rho}{\pi} L_{\text{pixel}} (1 - V(\omega_i)) \frac{A_{\text{s}} \cos\theta_{\text{s}_i} \cos\theta_{\text{r}_i}}{d_i^2}$$

Directional Occlusion

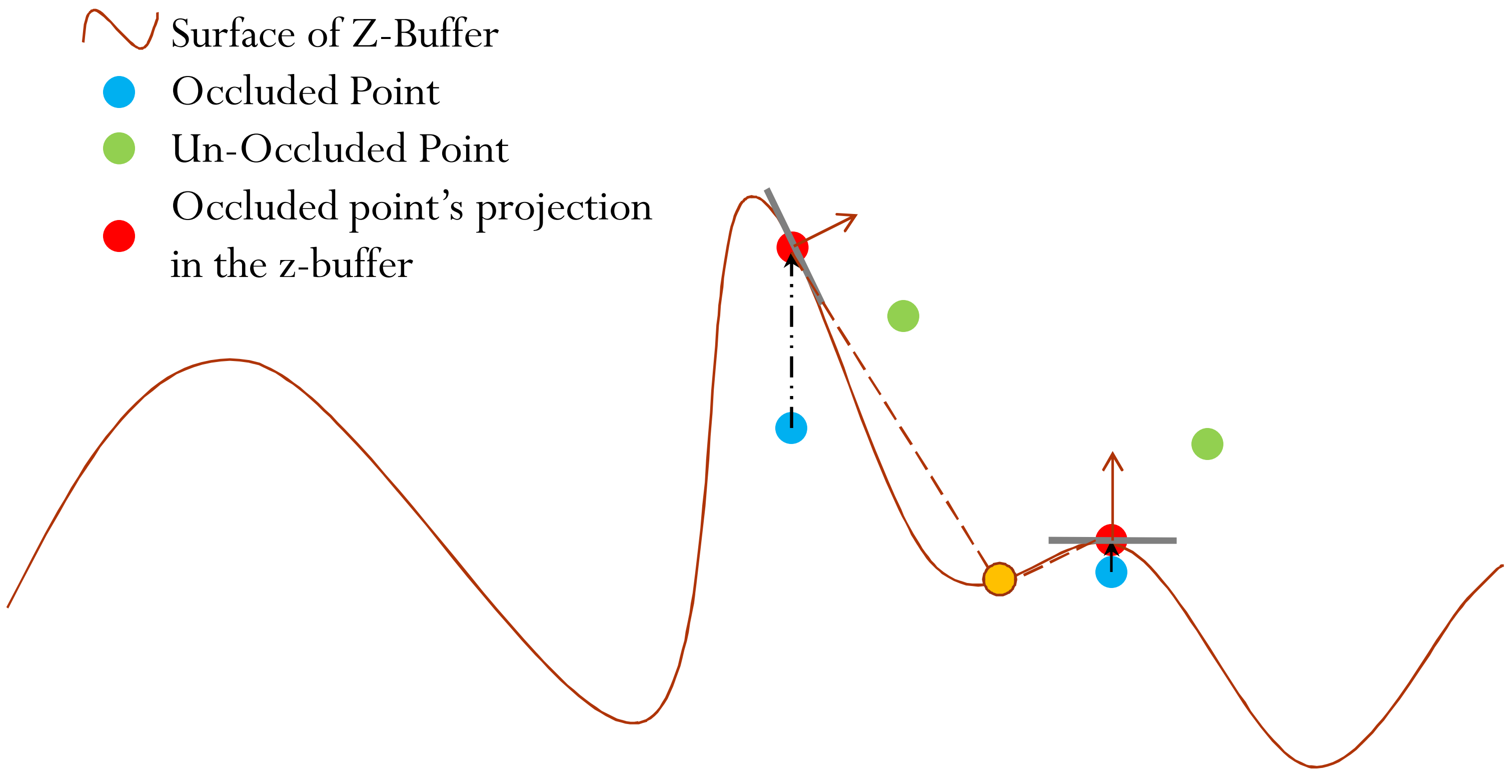

Comparison to a ray tracer

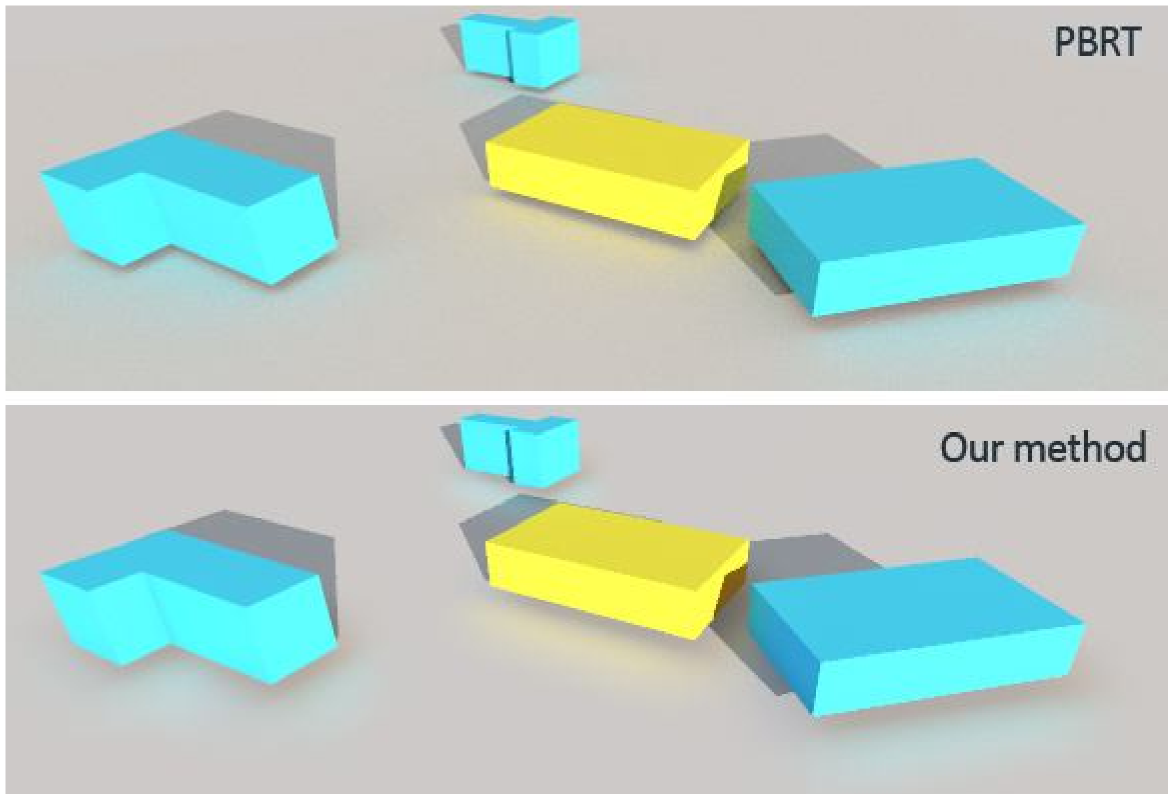

Deficiencies

Image space color bleeding does not work if the source surface is not visible.

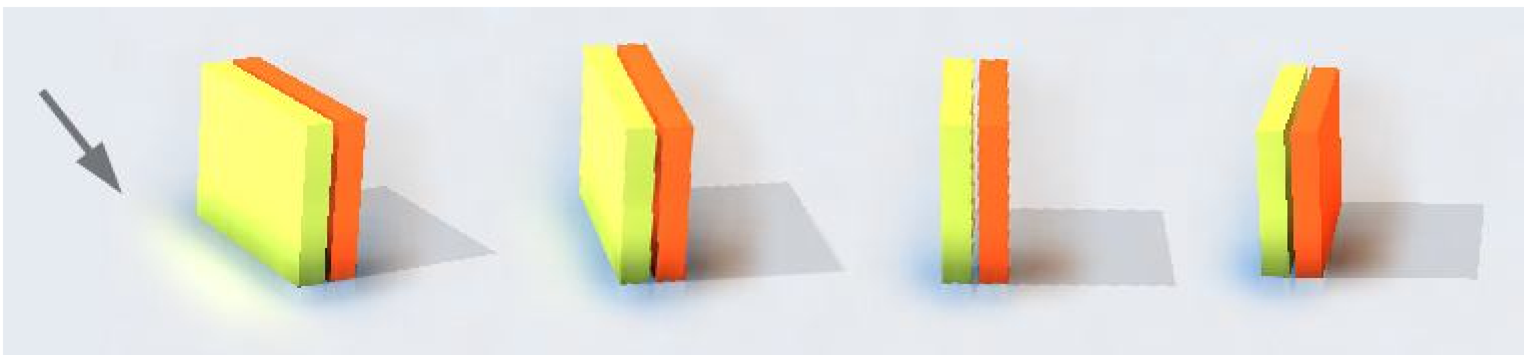

Directional occlusion also does not work when surfaces are not visible.

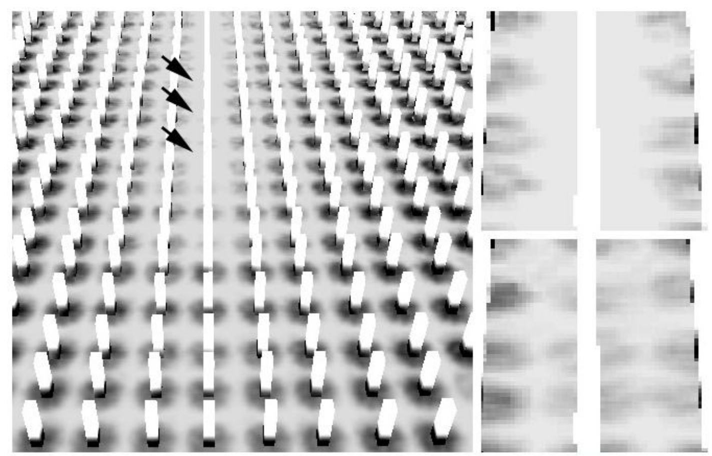

Problem with regular z-buffer

Each entry in the Z-Buffer stores the nearest depth value.

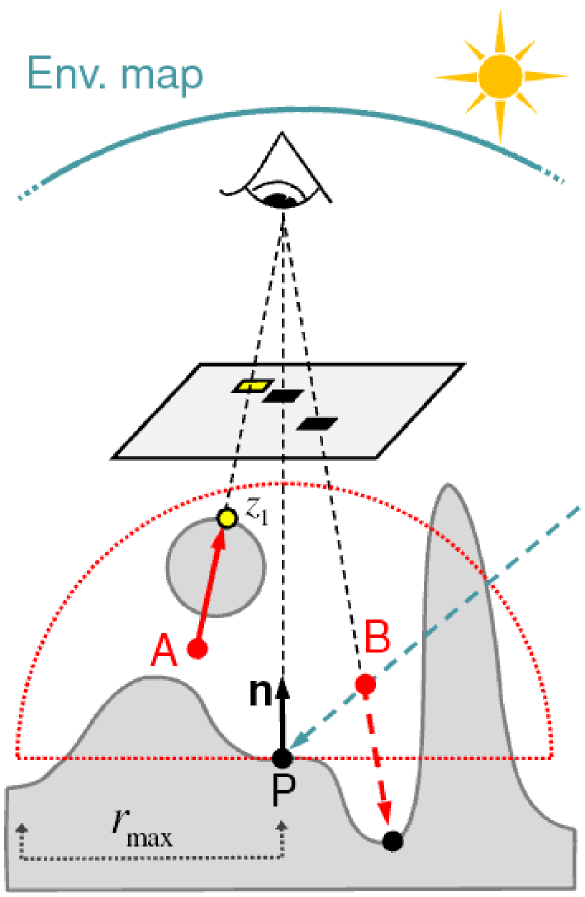

This means that point A is classified as an occluder, even though it is visible from point P.

Depth Peeling

Each entry in the Z-Buffer stores the ‘n’ nearest depth values.

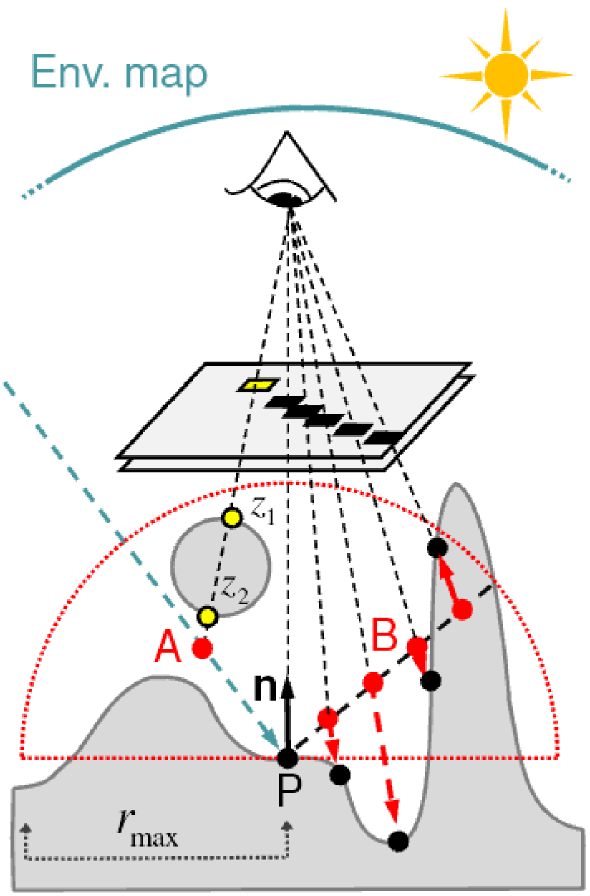

This means that point A is classified as visible, because it is below an even number of z values in the buffer.

Conclusions

- SSDO can deliver more realistic results than classic rasterization, with acceptable overhead (10% to 30% overhead reported as compared to SSAO).
- It is independent of object space complexity.
- There are no restrictions on dynamic scenes in terms of performance.

- Being an image space technique, it is still an approximation compared to physically based methods.

Credits for images

- The paper (Approximating Dynamic Global Illumination in Image Space)
- Wikipedia

Honor Code:

We pledge on our honor that we have not given or received any unauthorized assistance on this assignment or any previous homework.

QUESTIONS??