

# Non-Pinhole Approximations for Interactive Rendering




Figure 1. Non-pinhole depth image of teapot and reflections rendered with it (left), and conventional depth image and reflection rendered with it (right). The non-pinhole depth image captures the lid and bottom of the teapot and produces quality reflections.




Figure 2. Non-Pinhole relief texture maps and frames rendered with them (left) and conventional relief texture mapping (right).



Figure 3. Non-pinhole z-buffer illustrated with color (left), ambient occlusion effect rendered using it (middle), and ambient occlusion effect rendered using output image z-buffer (right). The non-pinhole z-buffer captures hidden parts of the dragon for a more complete and stable ground shadow.

## ABSTRACT

Depth images have been used to approximate scene geometry in a variety of interactive 3-D graphics applications. In previous work, images were constructed using orthographic or perspective projection which limits approximation quality to what is visible along a single view direction or from a single viewpoint. This paper shows that images constructed with non-pinhole cameras improve approximation quality at little additional cost provided that the non-pinhole camera offers fast projection. For such a camera, the fundamental operation of ray depth-image intersection proceeds efficiently by searching along the one-dimensional projection of the ray onto the image. In the context of two-camera configurations, our work extends epipolar geometry constraints to non-pinhole. We demonstrate the advantages of non-pinhole depth images in the context of reflections, refractions, relief texture mapping, and ambient occlusion.

**KEYWORDS:** non-pinhole camera, single-pole occlusion camera, graph camera, depth image, impostor, epipolar constraints, reflection, refraction, relief texture mapping, ambient occlusion, interactive 3-D graphics.

**INDEX TERMS:** I.3.3. [Computer Graphics]—Three-Dimensional Graphics and Realism.

## 1 INTRODUCTION

In the quest for higher-quality and higher-performance rendering researchers have resorted to approximating scene geometry with more efficient representations. The three main desirable properties of such an alternative representation are high-fidelity approximation, efficient construction, and efficient rendering. The representation should capture the geometry it replaces sufficiently well such that the resulting images be virtually indistinguishable from the images obtained when rendering with the original geometry. To support dynamic scenes, the representation has to be created on the fly, which requires fast construction.

Lastly, the alternative representation must deliver the desired performance boost to the application. We distinguish between applications where the representation can be rendered directly with the conventional feed-forward approach of projection followed by rasterization, such as for example when a distant tree is rendered using a billboard, and applications where the representation has to be rendered by intersection with one ray at a time, such as is needed for example in the case of reflections, refractions, relief texture mapping, and ambient occlusion. In this paper we focus on the second type of applications. For such applications a fast computation of the intersection between a ray and the alternative geometry representation is a central concern.

Images enhanced with per pixel depth have two of these desired properties. A depth image is constructed efficiently by rendering the geometry it replaces with the help of graphics hardware. Fast




Figure 4. Visualization of rays in 3-D (left) and of their curved projection on the non-pinhole depth image (right) used in Figure 1.

ray / depth image intersection is enabled by the fact that projection of a ray onto the depth image is a segment, which reduces the dimensionality of the intersection search space from two to one. Modern graphics hardware allows stepping along the ray projection, per pixel, at interactive rates. However, depth images are acquired from a single viewpoint—with a planar pinhole camera, or along a single view direction—with an orthographic camera, which limits their geometry modeling power. Such a depth image misses surfaces that become visible when the geometry approximation is rendered by the application, which lowers the quality of the result.

We propose constructing depth images using non-pinhole cameras. Such non-pinhole depth images offer a high-fidelity approximation of scene geometry while construction and rendering costs remain low. Once all rays need not pass through one point, the rays of a non-pinhole camera can be designed to sample all surfaces exposed by the application. To ensure construction and rendering efficiency, the non-pinhole camera is designed to provide a fast projection operation. This enables constructing the non-pinhole depth image in feed-forward fashion by projection followed by rasterization. The closed-form, unambiguous projection of the non-pinhole camera is leveraged a second time, during rendering, to compute the projection of the ray onto the non-pinhole image. As in the case of planar pinhole cameras, the ray / non-pinhole camera depth image intersection is found by walking on the one-dimensional projection of the ray. Unlike for planar pinhole cameras, ray projection is not a straight line, though this does not raise intersection costs significantly.

The advantages of non-pinhole depth images are demonstrated in the contexts of reflection, refraction, relief texture mapping, and ambient occlusion (also see accompanying video). In Figure 1 the non-pinhole depth image captures all teapot surfaces exposed in the reflections, whereas a conventional depth image produces incomplete reflections. The intersection between a reflected ray and the non-pinhole depth image is found by searching along the curved projection of the ray (Figure 4). Refractions are supported similarly by intersecting the non-pinhole depth image with refracted rays (Figure 5). Non-pinhole cameras greatly enhance the modeling power of relief texture mapping (Figure 2). A single-layer non-pinhole relief texture captures the entire top and sides of a barrel or an entire car to produce quality frames with rich detail. By contrast a single-layer conventional relief texture misses a considerable part of the barrel. A multilayered relief texture is also impractical in this case since a prohibitively large




Figure 5. Refractions rendered with a non-pinhole depth image of the teapot similar to the one shown in Figure 1.




Figure 6. Graph camera depth image (top) capturing an entire 3-D maze (bottom left) and reflection rendered using it (bottom right).

number of layers are needed to sample the side of the barrel well.

A promising technique for rendering with ambient occlusion at interactive rates is to use the output image z-buffer to approximate the exposure of output image samples to the environment [BS08]. However, the amount of occlusion for visible samples also depends on samples that are not visible in the output image. This makes the ambient occlusion effect unstable, with dark regions appearing and disappearing as geometry subsets appear and disappear in the output image. We propose using a non-pinhole z-buffer which captures most samples needed to estimate the occlusion of output image samples, producing a more complete and more stable ambient occlusion effect (Figure 3).

Abandoning the pinhole constraint presents the opportunity to design the camera model such as to obtain the depth image best suited for the application and dataset at hand. In this work we construct depth images by adapting two recently introduced non-pinhole camera models: the *single-pole occlusion camera* (SPOC) [MPS05], and the *graph camera* [PRA09]. The SPOC reaches around an object’s silhouette to gather samples not visible from the reference viewpoint but close to the silhouette. Such “barely” occluded samples are needed when the depth image is sampled from nearby viewpoints. The SPOC is well suited for approximating single objects, and it was used in the examples in Figures 1, 2, 4, and 5. The graph camera is a non-pinhole camera constructed starting from a planar pinhole camera through a series of frustum bending, splitting, and merging operations. The result is literally a graph of planar pinhole cameras. The graph camera circumvents occluders to sample an entire 3-D scene in a single-layer image. In Figure 6 the graph camera depth image models the entire reflected scene.

In summary, the contributions of our paper are as follows:

- We extend epipolar like constraints to non-pinholes. This fundamental contribution opens up the class of non-pinhole geometry approximations to interactive rendering applications.
- We provide a generic algorithm for intersecting a non-pinhole depth image with a ray, and we specialize the algorithm for two non-pinhole cameras, the SPOC and the graph camera.
- We demonstrate the advantages of non-pinhole depth images in the context of reflections, refractions, relief texture mapping, and ambient occlusion.

## 2 PRIOR WORK

We review prior research on image-based geometry approximation, on acceleration of rendering effects using depth images, and on non-pinhole cameras models.

## 2.1 Image-based geometry approximation

Maciel and Shirley [MS95] coined the term impostor which is now widely used to denote an image-based simplified representation of geometry for the purpose of rendering efficiency. The simplest impostor is a billboard, a quad texture mapped with the image of the original geometry, with transparent background pixels. Billboards are constructed efficiently, intersecting a billboard with a ray is trivial, and billboards provide good approximations of geometry seen orthogonally from a distance. When the billboard is close to the viewer the drastic approximation of geometry is unacceptable. Billboard clouds [DDS\*03] use several quads to improve modeling quality. The quads and the assignment to original geometry are optimized for modeling fidelity. The number of quads is small and a billboard cloud can be intersected with a ray one quad at a time. However, the optimization makes construction of the billboard cloud a lengthy process that precludes dynamic scenes. Moreover, the approximation quality is still insufficient for close-up viewing.

Depth images [MB95] greatly improve the modeling power of billboards. Constructing a depth image is just as inexpensive as constructing a billboard, but the cost of intersection with a ray is not constant anymore, but rather linear in the depth image width. Searching for the intersection in the entire image is avoided by leveraging epipolar-like constraints: the intersection is known to belong to the image plane projection of the ray. Since the depth image is constructed with an orthographic or a perspective projection, the depth image only captures samples visible along the reference direction or from the reference viewpoint. When surfaces not captured by the depth image are exposed by the application, objectionable disocclusion artifacts occur.

The simplest method for alleviating disocclusion errors is the use of additional depth images [MMB97], which is expensive and only palliative. A breakthrough came with the introduction of layered representations such as the multi-layered z-buffer [MO95] and the layered depth image (LDI) [Sha98], which allow for more than one sample along a ray and control disocclusion errors effectively. However, expensive construction restricts layered representations to static scenes. Moreover, the lack of a connected representation makes ray intersection difficult, precluding rendering effects such as those addressed by this paper.

As graphics hardware has evolved, it became possible to intersect depth images with individual rays, and research focus shifted to rendering effects involving higher order rays.

## 2.2 Rendering effects accelerated using depth images

Reflection and refraction have been studied extensively in interactive rendering, yet no complete solution exists. We assign reflection and refraction rendering techniques to four groups: ray tracing [Whi80], image-based rendering [LH96, GGS\*96, DYB98], projection [OR98], and reflected/refracted scene approximation. We only discuss group 4, since it is most relevant.

Environment mapping [BN76] approximates the reflected scene with cube map and it is currently the preferred approach for interactive applications due to its efficiency, robustness, and good results when scene geometry is far from the reflector/refractor. Environment mapping performs poorly close to the reflector/refractor. Improved results are obtained by approximating the scene with a sphere [Bjo04], but few environments are spherical so the fidelity is still quite limited. The scene approximation can be improved using depth images [SALP05, PMDS06]. Quality reflections are produced for simple objects or for select viewpoints, but the insufficient coverage is a limitation for non-trivial scenes or wide viewpoint translations.

Compared to reflection, refraction rays require additional work since most rays interact with the refractor at least twice—once

entering and once leaving the object. Several techniques have been developed for computing the second refraction at interactive rates, including pre-computed distance fields [CW05], GPU ray tracing [RAH07], and image-space approximations [Wym05].

Another rendering effect that requires intersecting depth images with individual rays is relief texture mapping [POC05]. True geometric detail is added to a coarse model by texturing each triangle with a height map. A conventional relief texture samples surface detail orthographically, along the direction of the normal of the underlying coarse model, which limits the technique to height field surfaces. Even so, sampling degrades when the geometric detail becomes aligned with the normal of the underlying surface. The technique has been extended to non-height field surface detail by resorting to a relief texture with multiple layers, each sampled orthographically [PO06]. The extension works well when complex detail can be captured in a few layers, as it is the case for a chain link fence for example. A strength of the extended method is the ability to capture double-sided detail. However, capturing geometric detail perpendicular to the underlying surface remains challenging as a large number of layers is needed. Our work extends relief texture mapping in an orthogonal direction and multilayered non-pinhole relief textures could be developed to collect the advantages of both techniques.

Ambient occlusion techniques add realism to local illumination models by approximating the amount of light a surface point receives based on how much of the environment is hidden from the point by nearby geometry. The computational cost is high since a ray has to be cast from each point in all directions. First implementations precomputed ambient occlusion in model space off-line [ZIK98], which precludes dynamic scenes. Initial efforts of enrolling the GPU in accelerating ambient occlusion resorted to a large number (i.e. 128-1024) of spherical shadow maps of the scene [PG04]. Our non-pinhole z-buffer ambient occlusion method builds upon an image-space technique introduced by Bavoi and Sainz [BS08]. Their technique approximates the amount of ambient occlusion at an output pixel using the output image z-buffer. They noticed that in order to sample occlusion at a pixel for an entire half plane it is sufficient to traverse one z-buffer segment. The result is a fast ambient occlusion method that supports dynamic scenes. However, the technique computes ambient occlusion as if the geometry seen by the output image were the only geometry in the scene, which can cause missing and unstable ambient occlusion artifacts.

## 2.2 Non-pinhole cameras

Non-pinhole cameras have been studied relatively little in computer graphics. The light field [LH96] and the lumigraph [GGs\*96] can be seen as the color samples acquired by a 2D array of planar pinhole cameras. Their strengths lie in the acquisition of small-scale complex real-world scenes. Although possible in principle, using light fields as geometry approximations is precluded by their large memory footprint and construction time. Multiple-center of projection cameras [RB98] sample the scene with a vertical slit along a user chosen path and thus avoid redundancy and offer sampling flexibility. However, construction requires rendering the scene for each position along the path. Camera models developed for multi-perspective rendering simulate camera motion through a 3-D scene but do not support novel views or dynamic scenes [Woo97, YM04].

Occlusion cameras address disocclusion errors. Given a reference view and a 3-D scene, an occlusion camera builds a single-layer image that stores not only samples visible from the reference viewpoint, but also samples visible from nearby points. In addition to the single-pole occlusion camera (SPOC) discussed earlier, other occlusion cameras include the depth discontinuity occlusion camera (DDOC) [PA06] and the epipolar occlusion

camera (EOC) [RP08]. Whereas the SPOC specifies the 3-D distortion of the reference view rays analytically, the DDOC specifies the distortion through a map. The added flexibility comes at the cost of increased construction times. The EOC captures all samples visible as the viewpoint translates between two given points. The EOC effectively generalizes the viewpoint of a planar pinhole camera to a viewsegment. However, the EOC only supports translation along a single direction.

### 3 NON-PINHOLE CAMERA DEPTH IMAGES

After removing the pinhole restriction, there is great flexibility in devising a camera model that best suits a given application and a dataset. Therefore we first discuss construction and ray intersection for non-pinhole depth images in general.

#### 3.1 Construction

Given a non-pinhole camera with a fast projection operation that maps a 3-D point  $(x, y, z)$  to  $(u, v, z)$  where  $(u, v)$  are image coordinates and  $z$  is a measure of depth linear in image space, a non-pinhole depth image is constructed efficiently by projecting the vertices of the geometry it is called to approximate and by rasterizing the projected triangles conventionally. The unconventional projection occurs in a vertex program which implements the non-pinhole camera model. Since lines do not necessarily project to lines anymore and since rasterization parameters do not vary linearly (before the perspective divide) anymore, the triangles must be sufficiently small to provide an adequate approximation. Complex objects are frequently modeled with small triangles, so additional tessellation is usually not needed. Meshes with large triangles can be subdivided on-the-fly by taking advantage of primitive-level GPU programmability.

#### 3.2 Ray intersection

Like a regular depth image, a non-pinhole depth image is defined by an image with color and depth per pixel, and a camera model which allows projection. The intersection of a ray  $(a, b)$  with a non-pinhole depth image  $NPI$  is computed as follows:

1. Clip the segment  $(a, b)$  with the bounding volume of  $NPI$  to obtain the segment  $(c, d)$ , see Figure 7.
2. Interpolate  $(c, d)$  in 3-D, from near to far to create  $n$  sub-segments. For each sub-segment  $(s_k, s_{k+1})$ :
  - 2.1. Project  $s_k$  and  $s_{k+1}$  to depth image at  $p_k = (u_k, v_k, z_k)$  and  $p_{k+1} = (u_{k+1}, v_{k+1}, z_{k+1})$ .
  - 2.2. Lookup image depths  $iz_k, iz_{k+1}$  at  $(u_k, v_k), (u_{k+1}, v_{k+1})$ .
  - 2.3. Intersect in 2-D segments  $[(0, z_k), (1, z_{k+1})]$  and  $[(0, iz_k), (1, iz_{k+1})]$  to obtain intersection  $(t_j, z_j)$ .
  - 2.4. If  $(t_j, z_j)$  is a valid intersection, return depth image color  $ic_j$  at  $\text{lerp}((u_k, v_k), (u_{k+1}, v_{k+1}), t_j)$ , else continue with next sub-segment.

The ray is interpolated in 3-D since its projection is not a straight line, and one cannot simply rasterize the segment that connects the endpoint projections. Each intermediate point is projected with the non-pinhole camera to trace the curved projection correctly. Since the depth  $z$  stored by the depth image varies linearly in the image, the intersection can be computed efficiently in a 2-D space  $(t, z)$ , where  $t$  is the parameter locating the intersection along segment  $(p_k, p_{k+1})$ . These generic algorithms are specialized for the SPOC and the graph camera as follows.

### 4 SINGLE-POLE OCCLUSION CAMERA DEPTH IMAGES

The SPOC projection uses a conventional projection followed by a distortion that moves the projected sample away from a pole

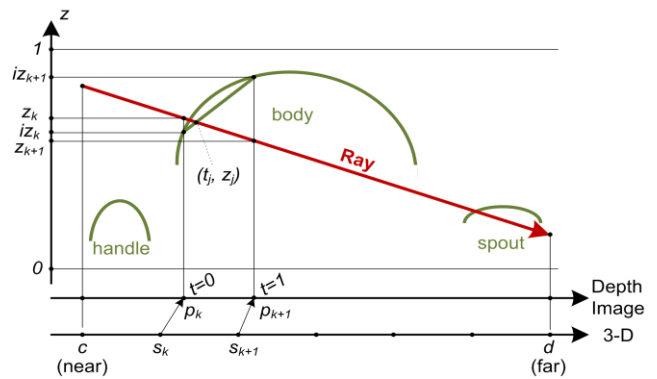


Figure 7. Ray / non-pinhole depth image intersection.

[MPS05]. The pole is the projection of the center of the object. The distortion magnitude increases with depth, so deeper samples move more, escaping the occluding front surface. For the SPOC depth image in Figures 1 and 2 the distortion pushes the object silhouettes back, revealing the lid and the bottom of the teapot, the entire side of the barrel, and the side and wheels of the car. Figure 8 shows that the SPOC depth image captures about half of the teapot, which is sufficient to intercept all reflected rays that would intersect the original teapot geometry.

SPOC construction and intersection closely follow the algorithms described in the previous section. The number of sub-segments  $n$  is chosen as the Euclidian distance between the projection of the endpoints of the clipped ray. This provides a good approximation of the actual number of pixels covered by the curved projection of the ray. The curved projections of a set of coplanar rays are visualized in Figure 4.

### 5 GRAPH CAMERA DEPTH IMAGES

The graph camera is constructed from a planar pinhole camera through a succession of bending, splitting, and merging operations [PRA09]. The result is a graph of planar pinhole camera frusta. The concept of camera ray is generalized to the set of points projecting at a given image location, which allows for rays that are not straight lines. The rays of the graph camera are piecewise linear. A ray changes direction as it crosses the shared face separating two frusta, but it remains continuous, which makes the graph camera image continuous. The graph camera constructed for the maze in Figure 6 is shown in Figure 9. The construction used a breadth first traversal starting from the entrance.

Projecting a point with the graph camera uses two steps. First, the frustum containing the given 3-D point is found. Then the point is projected directly to the output image with a 4-D matrix that concatenates the projections of all the cameras on the path from the current frustum to the root. The frustum containing the point can be found with an octree or another hierarchical space subdivision [PRA09]. Using the projection, graph camera depth image construction proceeds as in Section 3.1, except that a triangle has to be processed with each frustum it intersects.




Figure 8. Visualization of samples stored by conventional (left) and SPOC (right) depth image.




Figure 9. Graph camera model visualization. The frusta are shown in red and few rays are shown in white.

### 5.1 Ray intersection

In order to intersect a graph camera depth image with a ray it is possible to follow the generic algorithm: interpolate the ray uniformly in 3-D space, and project each new point onto the graph camera image. However, since each frustum is a pinhole, the projection of a given ray is piecewise linear (Figure 10), which enables the following optimization. Given a ray  $r$ , for each graph camera frustum  $F_i$ :

1. Intersect ray  $r$  with  $F_i$  to produce 3-D sub-segment  $(s_i, e_i)$ .
2. Project segment  $(s_i, e_i)$  to graph camera image segment  $(p_i, q_i)$ .
3. Walk on  $(p_i, q_i)$  to find intersection.

The algorithm computes the linear pieces of the ray directly by intersecting the ray with all the frusta, resulting in a set of sub-segments  $(s_i, e_i)$ . This is more efficient than the generic algorithm which requires small 3-D steps just to model the breaking points of the piecewise linear with high fidelity. Each frustum is a planar pinhole camera, so each sub-segment projection remains a straight line segment  $(p_i, q_i)$  in the output graph camera image. The sub-segment is interpolated to search for the intersection step by step, similarly to the generic algorithm.

## 6 APPLICATIONS TO INTERACTIVE RENDERING

Non-pinhole depth images accelerate reflection, refraction, relief texture mapping, and ambient occlusion as follows.

### 6.1 Reflection and refraction

To render a frame of a scene with specular reflections the first step is to update the depth images approximating reflected geometry that is dynamic. Then each reflector is rendered by computing a reflected ray per pixel, a step similar to environment mapping, and by intersecting the reflected ray with the depth images of reflected geometry. Multiple reflections of the same object are obtained at no extra cost (Figure 11, left). Fully dynamic scenes are supported since our method does not require any pre-computation involving the reflector and the non-pinhole depth images are computed efficiently. Higher-order reflections




Figure 10. Visualization of 3-D ray (left) and its piecewise linear projection in a graph camera image (right).



Figure 11. Multiple (left) and 2<sup>nd</sup> order (right) reflections.

are supported by storing per pixel normals. The normal at a first intersection point is used to create a second order reflected ray which intersects depth images again (Figure 11, right).

When deciding how to approximate reflected geometry the goal is to devise the simplest approximation that captures all samples visible in reflections. For example the black and white ground plane in the reflections in Figure 11 is perfectly captured with a billboard. The SPOC depth image of the teapot is captured along the direction that connects the centers of the bunny and teapot. If the teapot were spinning or if there were multiple reflectors surrounding the teapot, the best solution would be to use two depth images capturing complementary halves of the teapot.

Refractions are rendered by intersecting the emerging ray with the depth images approximating geometry. The algorithm for computing emerging refracted rays is orthogonal to this work. We use an image-space approximation for computing the emerging refracted rays [Wym05]. The key idea behind this approximation is to use a first rendering pass to store depth and surface normals for back-facing surfaces, which are then used by a second pass to compute the ray emerging after a second refraction.

### 6.2 Relief texture mapping

Relief texture map rendering is triggered by rendering the primitives of the coarse underlying model. To obtain correct silhouettes we render the bounding box of each relief tile (Figure 12). For every pixel the eye ray is transformed to the coordinate system of the current relief tile and intersection proceeds as before. World space  $z$  is computed at the intersection for correct  $z$ -buffering with the rest of the scene and for casting and receiving correct shadows. Shadows could be computed by shooting a second ray from the intersection to the light source and intersecting it with the relief texture. We prefer to use a conventional shadow map such that the relief surface casts and receives shadows from other objects and from other relief tiles.

Non-pinhole relief textures capture complex objects in a single layer. Figure 13 shows that a conventional relief texture misses the wheels and severely under-samples the sides of the car. With an increased complexity of the geometry modeled with the relief texture comes the desire to modulate the appearance of individual instances of the relief texture. The different colored cars in Figure




Figure 12. SPOC relief texture (left) and output image with relief bounding box visualization (right).




Figure 13. Conventional relief texture (left) & output image (right).  
 2 were obtained with a single relief texture by simply modifying the color of the intersection if it corresponds to the car body, identified through its yellow color.

### 6.3 Ambient occlusion

The horizon-based screen-space ambient occlusion algorithm [BS08] is fast because one can estimate occlusion in a half plane by traversing a single output image z-buffer segment. In Figure 14 left,  $p$  is an arbitrary output image pixel. Occlusion on the  $q$  side of  $ep$  is estimated by simply traversing  $pq$ , and casting rays  $r_i$  is *not* needed. This is enabled by the property that at any output pixel  $p$  any plane through the pixel ray  $ep$  will project to a line. This property needs to be maintained when porting the algorithm to non-pinhole z-buffers. The SPOC does not have this property. For any pixel other than the pole there is no set of planes spanning space with each plane projecting to a curve. The planes have a projection with non-zero area. We have designed a graph camera model that enhances the output image z-buffer with samples visible from two nearby viewpoints and also exhibits the desired property. The graph camera projection is equivalent to a series of conventional projections. Once a plane is collapsed to a line by the leaf projection, the line is mapped to lines by subsequent projections and the property is maintained. Using Figure 14 left again,  $pq$  remains a line after subsequent projections.

Figure 14 right visualizes the rays of the graph camera used to render the non-pinhole z-buffer shown in Figure 3. The graph camera has 3 sub-frusta: the output image frustum  $e$  up to a vertical plane through the splitting point  $s$ , and the left and right frusta  $l$  and  $r$  beyond. The splitting point moves on the blue curve as the output image view revolves around the dragon. The curve was designed off-line to move the splitting point smoothly behind the dragon as the dragon is seen sideways, in which case a conventional z-buffer suffices (also see accompanying video).

## 7 RESULTS AND DISCUSSION

### 7.1 Quality

Non-pinhole depth images enable quality reflections, refractions, relief texture mapping, and ambient occlusion as attested by the




Figure 14. Screen-space ambient occlusion algorithm (left) and visualization of graph camera model (right) used in Figure 3.




Figure 15. Comparison of our method (left) to ray tracing (right).

images in the paper and by the accompanying video. Figure 15 shows that our method achieves results comparable to ray tracing. The main limitations of our approach are as follows.

*Absent self-reflection.* Although our method could, in principle, support self-reflections by also intersecting the reflected rays with a depth image of the reflector, the additional intersection is probably a price interactive applications are not willing to pay.

*Coarse silhouettes.* An SPOC depth image does not sample the entire object it replaces. The sampled area ends with a jagged edge when the SPOC rays are tangential to the replaced geometry (Figure 8). When the jagged edge is exposed, the silhouette of the reflection becomes coarse. One possible solution is to smooth the edge as a pre-process, an approach that precludes dynamic scenes. Instead we alleviate the problem at run time by alpha blending the intersection sample with greater transparency when the SPOC ray becomes tangential to the sampled surface.

*Undersampling.* Like for all sample-based methods, the quality of the results obtained with non-pinhole depth images is contingent upon adequate sampling. The SPOC sampling rate is uniform and controllable. The graph camera sampling rate is not uniform: it is higher closer to the initial frustum and is lower for the distant frusta. The graph camera depth image used here was constructed to capture the entrance at a higher resolution, where reflections are of highest quality (Figure 6). Deeper into the maze the resolution decreases leading to aliasing artifacts due to the large output image projection of depth image pixels, a problem similar to inadequate shadow map resolution. In Figure 16 the teapot is in the top left corner of the maze (see Figure 6), thus deepest in the graph camera depth image, where sampling insufficiency is noticeable. Note however that the case presented here is particularly challenging: a smooth highly-specular surface reflects a contrasting checker pattern. We use a graph camera depth image resolution of 1,920 x 1,175. A brute force solution is to increase the resolution further. Another possibility is to break up the maze into several parts each with its own smaller graph camera depth image. Lastly, hybrid sample-based/geometry-based




Figure 16. Undersampling artifacts on floor reflection.




Figure 17. Diffuse teapot reflected in body of large teapot (left) and visualization of number of intersection steps per pixel (right).




Figure 18. Silhouette detail with fine steps of 1,  $\frac{1}{2}$ , and  $\frac{1}{4}$  pixels.

techniques that incorporate “infinite frequency” edges into textures could also be used.

*Missing samples.* The most visible artifacts in non-pinhole relief texture mapping are caused by samples still missing from the relief texture due to residual occlusions. The rear bumper of the car shown in Figure 2 occludes some of the car body in the relief texture, which causes the shimmering “rubber band” surface seen in the video. One solution is to modify the car model to reduce the distance between the bumper and the body of the car by pushing the bumper in or by thickening the bumper. Another solution is to encode the bumper in a second relief texture layer.

## 7.2 Performance

The timing information reported here was collected on a 3.4GHz 2GB Intel Xeon PC with an NVIDIA 8800 Ultra 768MB card. We used NVIDIA’s Cg 2.0 shading language with gp4 profiles. An important performance factor is the number of steps along the projection of the ray, which we analyzed for reflections.

We take coarse steps first and perform a fuzzy intersection of the coarse ray segment with the non-pinhole depth map. If the two endpoints project at unoccupied locations or if the coarse ray segment clearly does not intersect the impostor depth map, the coarse segment is trivially rejected. Coarse segments are refined by performing fine steps of 1,  $\frac{1}{2}$ , or  $\frac{1}{4}$  depth image pixels. Figure 17 illustrates the number of steps for a 512x512 SPOC depth image, a 6 pixel coarse step, and a  $\frac{1}{4}$  pixel fine step. More steps are needed when the reflected ray narrowly misses the teapot, which causes the fuzzy test to return a false positive. The average number of steps is 48 per output pixel, including both coarse and fine steps. For fine steps of 1 and  $\frac{1}{2}$  pixels the average number of steps is 22 and 31, respectively. These numbers do not account for pixel processor idling due to SIMD processing constraints. Figure 18 shows reflection silhouette quality for various fine step sizes.

Performance depends on output image resolution and on the fine step size as shown in Table 1. Performance was measured on a typical path (see video) for the scene shown in Figure 1. Eight

Res.	640x480			800x600			1024x768		
<b>Fine step</b>	<b>1</b>	<b>1/2</b>	<b>1/4</b>	<b>1</b>	<b>1/2</b>	<b>1/4</b>	<b>1</b>	<b>1/2</b>	<b>1/4</b>
<b>Avg FPS</b>	26.5	20.7	15	20.9	16.5	11.6	15.2	11.8	8
<b>Min FPS</b>	18	14	8	14	10	4	10	8	6
<b>Max FPS</b>	54	46	36	48	15	34	40	34	28

Table 1. Frame rate [fps] variation with output resolution and fine step size [pixels] for scene in Figure 1.



Figure 19. Teapot location and reflection (top), and number of steps visualization (bottom) w/o and w/ coarse stepping. Brighter red indicates a large number of steps being taken.

sample multi-sampling anti-aliasing (8x MSAA), a 512x512 SPOC depth image, and a coarse step of 6 pixels were used. For an output resolution of 640x480, with MSAAx8, the average frame rate for SPOC depth images of resolution 128x128, 256x256, 512x512, and 1024x1024, is 55.8, 36.6, 26.5, and 16.1fps, respectively. For coarse steps of 3, 6, 9, and 12 texels, the average frame rates are 18.2, 31, 37.2, and 39.8fps, respectively. The only feature thin enough to be affected by the coarser steps is the tip of the spout. For a sequence where the SPOC depth image is recomputed on the fly (see video), the average frame rates are 22 and 17.3fps for no anti-aliasing and 16x MSAA, respectively.

Coarse stepping reduces the number of steps for the graph camera depth image as well, as can be seen in Figure 19, where the average number of steps decreases from over 155 to 27. The reflection of the main entrance where the graph camera impostor has highest resolution remains the hottest area on the teapot but only a few pixels have large step numbers. The graph camera impostor is constructed at over 100fps. The average, min, and max frame rates for the path that follows the teapot through the maze (see video) are 45.5, 30, and 105fps w/o antialiasing, and 26.8, 20, and 42fps with 8x MSAA, respectively.

We will examine accelerating the ray/depth image intersection computation further by leveraging ray coherence. We envision a two pass approach that first renders the reflection at lower resolution and then up-samples by interpolation in coherent regions. The second pass actually computes intersections only at regions where the lower resolution results are not sufficient to reconstruct a quality intersection, such as at edge regions.

We construct non-pinhole relief texture maps with an SPOC and the discussion of performance of ray intersection provided above in the context of reflections still applies. For Figure 2 the overall performance, including shadow mapping, is 14fps for the 40 cars and 18fps for the 60 barrels example. Output resolution is 640x480 and relief texture resolution is 512x512. For 20, 10, and 1 car performance is 26, 51, and 219fps, respectively. All examples shown use tall relief which implies long ray projections. For scenes with short relief performance is even higher. For example, for 160 cars half the size (Figure 20) performance is 46fps.

We investigated ambient occlusion performance for two quality occlusion sampling settings: regular (6 sampling directions and 6 steps per direction), and fine (32 and 20). The blur




Figure 20. Short relief example.

kernel width was 21 pixels and the output image resolution was 1024x1024 in both cases. Average performance for the dragon scene (Figure 3) is 35 and 16fps for the regular and fine settings, respectively. Figure 3 and the video use the fine sampling setting—the regular setting produces a noisier ambient occlusion, issue orthogonal to the use of the non-pinhole z-buffer. Rendering the locally illuminated output image, the graph camera z-buffer, and adding the ambient occlusion effect takes 5.9, 16, and 6ms for the regular setting and 5.9, 16, and 40ms for the fine setting, respectively. For the conventional algorithm rendering the locally illuminated output image and adding the ambient occlusion effect takes 5.9 and 5.9ms for the regular setting and 5.9 and 32ms for the fine setting, respectively. Sampling occlusion in the graph camera z-buffer as opposed to the conventional z-buffer brings a 25% penalty, and most performance loss is brought by having to render the graph camera z-buffer of the 890K triangles scene. However, in some scenes it is unnecessary to update the graph camera z-buffer for every frame. The non-pinhole z-buffer does see more than what is visible in the output image which greatly extends its resiliency to output view changes. For example the z-buffer shown in Figure 3 can be reused if the dragon is only seen from the front, which increases performance to 21fps for the high setting, approaching the 26fps of the conventional algorithm. Another possible approach is to use a simplified version of the model when computing the non-pinhole z-buffer.

## 8 CONCLUSIONS AND FUTURE WORK

We have shown that the desirable property of conventional depth images of efficient ray intersection holds for non-pinhole depth images, at the small additional cost of traversing a curved—as opposed to a straight—ray projection. We have leveraged this property to render specular reflections, refractions, relief texture mapping, and ambient occlusion.

Compared to reflection rendering techniques that approximate the projection of reflected points such as explosion maps [OR98], our method produces multiple projections of the same object at no extra cost and handles complex reflectors. Compared to color-caching image-based techniques, our method supports dynamic scenes and has reduced memory requirements. Color-caching techniques excel at capturing the appearance of complex real-world materials that are glossy, but not specular. Compared to environment mapping, our method produces better results close to the reflector, at a higher per-pixel cost stemming from the more sophisticated scene approximation. Compared to ray tracing, our method more easily minifies and magnifies reflections by working in the color map at different levels of resolution, and achieves fast ray / geometry intersection. Ray tracing has a quality advantage since it does not approximate the reflected geometry. Compared to conventional relief texture mapping the non-pinhole relief maps bring greater modeling power at the cost of a costlier ray/relief texture intersection, but without the cost of additional relief texture layers.

In addition to the directions for future work already sketched, we will investigate integrating our non-pinhole rendering framework into popular digital 3-D content creation tools. Our work argues for the benefits and practicality of exploring camera models that abandon the pinhole constraint. Non-pinhole cameras can be designed to provide powerful yet inexpensive approximations for many applications in graphics and beyond.

## REFERENCES

- [BS08] BAVOIL L. AND SAINZ M. Image-Space Horizon-Based Ambient Occlusion, SIGGRAPH 2008, Talk Program, 2008.
- [Bjo04] BJORKE K. Image-based lighting. GPU Gems, Fernando R., (Ed.). NVidia, (2004), pp. 307–322.
- [BN76] BLINN J.F., NEWELL M. E. Texture and Reflection in Computer Generated Images. CACM 19:10, pp. 542-547, 1976.
- [CW05] CHAN B. AND WANG W. Geocube—GPU accelerated real-time rendering of transparency and translucency. The Visual Computer, vol 21, 2005, pp. 579-590.
- [DYB98] DEBEVEC P., YU Y., BORSHUKOV G. Efficient view-dependent image-based rendering with projective texture-mapping. In EG Workshop on Rendering, 1998, pp. 105–116.
- [DDS\*03] DECORET X. ET AL. Billboard Clouds for Extreme Model Simplification. ACM SIGGRAPH 2003, pp.689-696.
- [GGS\*96] GORTLER S., GRZESZCZUK R., SZELISKI R., AND COHEN M. The Lumigraph. In Proc. of SIGG. 96, pp.43-54.
- [LH96] LEVOY M., AND HANRAHAN P. Light Field Rendering. Proc. of SIGGRAPH 96, pp. 31-42.
- [MS95] MACIEL P., AND SHIRLEY P. Visual Navigation of Large Environments Using Textured Clusters. I3D'95, pp. 95-102.
- [MMB97] MARK W., MCMILLAN L., AND BISHOP G. Post-Rendering 3D Warping. Symp. on Interactive 3D Graphics, pp. 7-16.
- [MO95] MAX N. AND OHSAKI K. Rendering trees from precomputed z-buffer views. In Rendering Techniques '95: Proceedings of the EGRW, pp. 45–54.
- [MB95] MCMILLAN L. AND BISHOP G. Plenoptic modeling: An image-based rendering system. SIGGRAPH '95, pp. 39-46.
- [MPS05] MEI C., POPESCU V., AND SACKS E. The Occlusion Camera. Eurographics 2005, pp. 335-342.
- [OR98] OFEK E. AND RAPPOPORT A. Interactive reflections on curved objects. In Proc. of SIGGRAPH '98, ACM Press, 333-342.
- [PG04] PHARR M. AND GREEN S. Ambient Occlusion. In GPU Gems, edited by Randima Fernando, pp. 279–292, 2004.
- [PO06] POLICARPO F. AND OLIVEIRA M. Relief Mapping of Non-Height-Field Surface Details. In Proceedings of ACM Symposium on Interactive 3-D Graphics and Games 2006, pp. 55-62.
- [POC05] POLICARPO F., OLIVEIRA M., AND COMBA J. Real-Time Relief Mapping on Arbitrary Polygonal Surfaces. In ACM Symp. on Interactive 3-D Graphics and Games 2005, pp. 155-162.
- [PRA09] POPESCU V., ROSEN P., AND ADAMO-VILLANI, N. The Graph Camera. ACM SIGGRAPH ASIA '09, to appear.
- [PMDS06] POPESCU V., MEI C., DAUBLE J., AND SACKS E. Reflected-Scene Impostors for Realistic Reflections at Interactive Rates. Computer Graphics Forum, (25):3 (EG 2006), pp. 313-322.
- [PA06] POPESCU, V. AND D. ALIAGA. The Depth Discontinuity Occlusion Camera. In ACM I3D 2006, pp. 139-143.
- [RB98] RADEMACHER P. AND BISHOP G. Multiple-center-of-Projection Images. Proc. ACM SIGG. '98, pp. 199–206.
- [RAH07] ROGER D., ASSARSSON U., AND HOLZSCHUCH N. Whitted Ray-Tracing for Dynamic Scenes Using a Ray-Space Hierarchy on the GPU. In EGSR, 2007, pp. 99-110.
- [RP08] ROSEN P. AND POPESCU V. The Epipolar Occlusion Camera, In ACM Symp. on Inter. 3-D Graphics 2008, pp. 115-122.
- [SALP05] SZIRMAY-KALOS L. ET AL. Approximate Ray-Tracing on the GPU with Distance Impostors. Computer Graphics Forum, 24(3), 2005. pp. 171-176.
- [Sha98] SHADE J. ET AL. Layered Depth Images. In Proceedings of SIGGRAPH 98, 231-242.
- [Whi80] WHITTED T. An improved illumination model for shaded display. Comm. of the ACM (1980), 23, 6, pp. 343-349.
- [Woo97] WOOD D. ET AL. Multiperspective Panoramas for Cel Animation. In Proc. of ACM SIGG. '97, pp. 243-250.
- [Wym05] WYMAN C. An Approximate Image-Space Approach for Interactive Refraction. ACM ToG, vol. 24, no 3, pp. 1050-1053.
- [YM04] YU J. AND MCMILLAN L. A Framework for Multiperspective Rendering. In Proceedings of Eurographics Symposium on Rendering (EGSR), 2004.
- [ZIK98] ZHUKOV S., IONES A., AND KRONIN G. An Ambient Light Illumination Model. In EGRW '98, pp. 45–56.