

Computational Highlight Holography

Christian Regg¹

Szymon Rusinkiewicz²

Wojciech Matusik³

Markus Gross^{1,3}

¹ETH Zürich

²Princeton University

³Disney Research Zürich




Figure 1: Computational highlight holography enables embedding three-dimensional models into reflective or transmissive surfaces. The user starts by placing a three-dimensional model with respect to the reflective surface, and specifying a light direction and the observer’s field of view (a). Our algorithm finds a set of points that best depicts the model’s surface (b). Our optimization packs surface fragments corresponding to these points on the hologram surface (c). We use a computer-controlled engraver to manufacture the final hologram surface. This surface reproduces three-dimensional images of the model with correct binocular and motion parallax (d).

Abstract

Computational highlight holography converts three-dimensional computer models into mechanical “holograms” fabricated on (specular) reflective or refractive materials. The surface consists of small grooves with patches of paraboloids or hyperboloids, each of which produces a highlight when illuminated by a directional light. Each highlight appears in different places for different view directions, with the correct binocular and motion parallax corresponding to a virtual 3D point position. Our computational pipeline begins with a 3D model and desired view position, samples the model to generate points that depict its features accurately, and computes a maximal set of non-overlapping patches to be embedded in the surface. We provide a preview of the hologram for the user, then fabricate the surface using a computer-controlled engraving machine. We show a variety of different fabricated holograms: reflective, transmissive, and holograms with color and proper shading. We also present extensions to stationary and animated 2D stippled images.

1 Introduction

Three-dimensional displays have long been a fascination of humans, and over the years we have discovered and studied a variety of technologies generate apparent 3D depictions of objects and scenes [Okoshi 1976]. Many cues contribute to our ability to perceive depth, but two of the most important are binocular parallax and motion parallax. Stereopsis, the perceived fusing of two slightly different images of a scene observed by the two eyes, leads to a strong sensation of depth. Motion parallax refers to the different relative motions of objects at different depths in the scene,

as an observer moves with respect to it. Ever since the effects of these cues have been formalized and understood, there has been continuous progress in the development of devices that utilize these principles and display 3D images. These, in turn, have been used in entertainment, advertising, tele-presence, visualization, security markers, and art. Devices and systems for stereoscopic display have also prompted recent interest within the computer graphics community [Perlin et al. 2000; Matusik and Pfister 2004; Jones et al. 2007; Fuchs et al. 2008].

In this work we introduce a process called computational highlight holography, which was inspired by an analog process (scratch holography) previously explored by hobbyists. We begin with an arbitrary 3D model and compute the shape of a specular surface, such that reflections from this surface cause an apparent depiction of points on the 3D object, with correct binocular and motion parallax. We use a computer-controlled engraving machine to fabricate the resulting surface.

Computational highlight holography is based on a few simple observations. First, it uses the fact that a paraboloidal mirror focuses parallel incident light rays to pass through a single point. Therefore, an observer sees a virtual 3D point at the focus of the mirror, from a range of viewpoints. An obvious problem with using this approach to depict a 3D scene is that only a single scene point can be correctly represented. However, it has been observed by hobbyists [Beaty 1995; Beaty 2003], artists [Garfield 1981], and researchers [Plummer and Gardner 1992; Abramson 2000; Eichler et al. 2003] that we can make a simple trade-off: by using segments of different paraboloidal surfaces embedded within a planar surface, we can represent more 3D points at the cost of limiting the range of viewer positions. Within that range, the observer sees a collection of highlights on the surface, with each highlight exhibiting the correct apparent motion to provide correct binocular and motion parallax.

Thus far, the scratch holography technique has not been explored a great deal: only a few practitioners have produced them by hand, using a sharp tool such as a compass to make scratches in a slab of plastic or metal. By moving to a computational approach and using a computer-controlled cutting machine, we obtain greater precision, can produce a greater number of 3D points, and improve the usefulness and practical applicability of the method.

In this paper, we begin by deriving the surface geometry that results in an image of a given 3D point (Section 3.1). Next, we describe how to determine a set of points depicting relevant features of an input 3D model (Section 3.2) and how to optimize for a good embedding of the surface patches corresponding to these points on a hologram surface (Section 3.3). The result is a 3D surface model for the mechanical hologram, which we fabricate using a computer-controlled engraving machine (Section 4.1). We take advantage of the fact that these devices have become affordable for consumers and are capable of manufacturing surfaces with extremely high precision. Among the results we present are not only traditional reflective holograms, but also transmissive (refractive) holograms and extensions to color, pure 2D images, and animations.

Our method has a few important benefits. First, it allows us to embed a 3D model in practically any existing reflective or transmissive surface. Second, the manufacturing cost is low, and could be made even lower by fabricating a die and then mass-producing a large number of copies by embossing or casting. Third, the whole process is simple, which allows hobbyists, artists, and designers produce custom scratch holograms easily and affordably. Finally, scratch holograms work under natural illumination, such as sunlight, and allow for variants such as transmissive and color models.

Taken together, these benefits suggest that computational highlight holography holds tremendous potential to be used in many different applications. For example, at a large scale it could be used in architecture to embed holograms in glass buildings. It could be used at a smaller scale in order to add corporate logos or product ads. Finally, the technique could be used for security markers.

2 Related Work

Traditional holography. Dennis Gabor was the first to formulate the principles of holography [1948], for which he received the 1971 Nobel Prize in Physics. However, the first holograms were recorded in concurrent work by Yuri Denisyuk [1962] and Emmett Leith and Juris Upatnieks [1962]. In the recording process, coherent light reflecting from a scene undergoes interference with a reference beam, and the resulting fringes are recorded on a photographic plate. In the reproduction step, light from the reference beam is diffracted by the fringes, recreating the original wavefront.

Traditional holograms are difficult to record, requiring sub-wavelength vibration control, and they can only be viewed with coherent light sources (e.g., lasers). Practical holography took a huge step forward with the invention of white light transmission and reflection rainbow holograms by Stephen Benton [1969]. These have full color and can be viewed under natural light, but sacrifice parallax in the vertical direction. Rainbow holograms can be mass-produced using embossing of micro-surface reliefs in plastic films. However, it is difficult to scale this technique to large sizes, and designing and recording such holograms remains expensive and not commonly accessible.

Spatially-multiplexed parallax displays. These displays emit spatially and directionally varying light, by combining a printed high-resolution pattern with a means of converting some of the spatial resolution into angular variation. The first parallax barrier displays, called parallax stereograms, were invented by Frederic Ives [1903] and supported only 2 views. They were later extended to multiple views and labeled as parallax panoramagrams [1928]. These displays all use vertical slits offset from the display surface to multiplex rays depending on the viewer’s position, thus providing parallax in one direction.

Almost concurrently, Gabriel Lippmann [1908] proposed using an array of spherical lenses to multiplex spatial information. Later,

in the 1940s, cylindrical lens (or lenticular) sheets were introduced to convert spatial patterns into directional light variation. Displays based on lenticular sheets provide parallax only in one direction, as opposed to those with spherical lenses, but they achieve higher spatial resolution. Lenticular displays have proven popular, being widely adopted for advertising and art, due to both the availability of large, high-quality, and inexpensive lenticular sheets, and the relative ease of the content creation process.

Scratch holography. The phenomenon of mechanical (or scratch) holograms has been reported by researchers [Plummer and Gardner 1992; Abramson 2000; Eichler et al. 2003], artists [Garfield 1981], and hobbyists [Beaty 1995; Beaty 2003]. However, there are even some reports of producing the first mechanical holograms in the mid-1930s. While some of the researchers analyzed the desired surface geometry, the process of generating these holograms was always analog: producing circular scratches, by hand, on a plastic or metal surface.

While our work was inspired by this method, it differs in fundamental ways. Most importantly, our design process results in an optimized 3D surface shape that has the proper geometry to result in apparent 3D point positions. We show that this requires paraboloidal arcs, cut into rectangular grooves in the surface. In contrast, previous work has always used hand-made, shallow 2D scratches of unknown and arbitrary cross-section. These produce highlight positions that are not strictly correct, resulting in only *approximately* correct parallax (i.e., 3D shape distortion) under viewer motion. In addition, intersections between the scratches degrade quality, which we avoid by making our grooves nonoverlapping.

In addition to using the correct surface geometry, we make the process completely automated and digital—from input 3D model to physical milled surface. This allows us to create mechanical holograms with many more points than any previous attempt. We also explore both reflective and transmissive surfaces (prior work has only examined the reflective case), and add color to the 3D points.

Custom reflection design. Our work is also related to research on generating specular surfaces with desired reflection. The goal of reflector design for luminaires is to compute a surface model that produces a desired light distribution, given a fixed local position of a point light source [Patow and Pueyo 2005; Patow et al. 2007]. The work by Weyrich et al. [2009] deals with a similar problem but for distant or directional light sources. They are able to produce reflective surfaces that produce reflections according to complex, user-specified stencils. While we are also generating surfaces with custom and user-specified reflection, the aim of our work is very different. We are attempting to produce surfaces that depict virtual 3D points, with correct binocular and motion parallax.

3 Method

We begin by describing the mathematical derivation for a surface fragment that produces a virtual image of a 3D point. We consider both the reflective and transmissive cases. Next, we explain our algorithm to pick a point set that adequately depicts the given 3D model. Finally, we show how we fill the hologram surface with these fragments.

3.1 Geometric Design of Grooves

3.1.1 Setup

We consider a single groove on the hologram surface, which focuses a directional incident light to produce an image¹ at the 3D

¹ In standard optics terminology, this will be a virtual image of the light-source if p is below the surface, and a real image if p is above the surface.

point p . As shown in Figure 2, we assume that the hologram surface consists of small offsets from the xz plane, and is illuminated by a directional light (i.e., parallel incident light rays) of direction $\mathbf{l} = (0, -1, 0)^T$. We consider a range of viewer positions centered around the point \mathbf{v} and extending in the $\pm x$ direction, with the angle α denoting the total horizontal field of view.




Figure 2: Setup for derivation of the reflective patch corresponding to one virtual point p . The viewer positions lie along a segment centered at \mathbf{v} , and parallel light rays are incident in direction \mathbf{l} .

Zooming in on the groove (Figure 3), we see that its extent must be a line in the plane (technically a rectangle, since it has some small vertical extent w), since it is defined by lines passing through p and the line of viewer positions. This is already a key difference from most scratch holography setups, which use curved scratches to compensate for the fact that they cannot control the scratch cross-sectional profile.

The center of the groove is found by casting a ray from the central view position \mathbf{v} through the sample point p , and intersecting it with the hologram plane:

$$\mathbf{i} = \mathbf{v} - \left(\frac{v_y}{p_y - v_y} \right) (\mathbf{p} - \mathbf{v}). \quad (1)$$

The length of the groove depends on the specified FOV: half of the x extent is

$$d_3 = d_2 \tan \frac{\alpha}{2}, \quad (2)$$

where $d_2 = \|\mathbf{p} - \mathbf{i}\|$. To simulate occlusion (i.e., disappearance of the point from certain view directions), we modify the field of view for each point, including the possibility of a different FOV for each side of the groove. The rectangular region is therefore spanned by

$$\mathbf{c}_1 = \begin{pmatrix} i_x - d_3 \\ 0 \\ i_z - \frac{w}{2} \end{pmatrix} \quad \text{and} \quad \mathbf{c}_2 = \begin{pmatrix} i_x + d_3 \\ 0 \\ i_z + \frac{w}{2} \end{pmatrix}. \quad (3)$$

3.1.2 Reflection

For reflective holograms, the necessary shape within the groove is a patch of a paraboloid, since this is the shape that reflects a bundle of parallel rays to pass through p . The family of surfaces $S_d : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by

$$S_d(x, z) = \frac{1}{4d} ((x - p_x)^2 + (z - p_z)^2) + (p_y - d) \quad (4)$$

consists of paraboloids of revolution that reflect light incident from $\mathbf{l} = (0, -1, 0)^T$ onto the point p . These paraboloids are concave (convex), as seen from above the surface, if the sign of p_y is positive (negative).




Figure 3: Detailed view of the rectangular groove for one point.

Since there are several possible paraboloids with focus p , we are free to select an optimal one. To do this, we impose two additional conditions:

1. Because the surface will be fabricated using a subtractive (en-graving) technology, we ensure that the paraboloid lies entirely below (or at) the surface within the rectangular region.
2. Because we wish to remove as little material as possible, we do not want to make the groove too deep. Therefore, we ensure that at least one point on the paraboloid is exactly at the surface.

Satisfying these conditions requires a different procedure depending on whether the paraboloid is concave or convex, which we have seen depends on whether the virtual point p is above or below the surface:

- In the concave case, the highest point will be one of the four corners of the rectangle. Therefore, we simply plug each of the four corners in turn into equation (4), solving the resulting quadratic for d . We select the one that results in a paraboloid with entirely negative y within the rectangle.
- In the convex case, we find the (unique) paraboloid that is tangent to the xz plane—it can be shown that the point of tangency is $(p_x, 0, p_z)^T$. If this lies within the rectangle, we plug it into (4), else we use the closest point to this on the boundary of the rectangle.

In either case, we are able to solve for a unique paraboloidal patch (Figure 4, top).




Figure 4: **Top:** Surface patch for a virtual point below the surface: this is a portion of a convex paraboloid. **Middle:** Surface patch after the addition of ramps. **Bottom:** Dividing the surface into several surface segments limits total depth.

Ramps. The above derivations may be generalized to consider different incident light directions, as well as viewer positions on a line that is not parallel to the plane. In this case, there is a danger that the paraboloidal patch is occluded by the edge of the groove, from the point of view of either the viewer or light source. To combat this, as well as aid in manufacturing and polishing, we give the edges of the groove a sloped ramp, as shown in Figure 4, middle.

Discontinuities. For certain settings (viewer and model position, field of view) some grooves are very deep. However, the production process allows only a limited plate thickness and deep grooves are difficult to polish, which limits the maximum depth grooves may have. A simple solution could be to cut off the quadric surface at a certain depth, which obviously leads to the absence of the desired effect at that location. An additional problem with deep grooves is that, already without manually adding ramps, the sides are tapered by the drill head because of its conical shape. This prevents us from being able to tightly pack grooves with large depth.

We solve these problems using a strategy inspired by Fresnel lenses: dividing the groove surface into smaller segments that are parts of different quadric surfaces having a certain maximum depth. Figure 4, bottom, shows a result of this process.

3.1.3 Refraction

It is also possible to construct a transmissive surface, such that the viewer looks through it at a directional light and again sees virtual 3D point positions. The ideal shape for such a surface is flat on the side seen by the light, and a hyperboloid of revolution on the surface seen by the viewer.

The equation for a circular hyperboloid of two sheets, with its major axis along y , is

$$\frac{y^2}{a^2} - \frac{x^2 + z^2}{b^2} = 1. \quad (5)$$

For a viewer position on the $+y$ side of the surface and light rays incident in direction $\mathbf{l} = (0, 1, 0)^T$, we use the lower sheet of the hyperboloid, as shown in Figure 5. In this case, the light rays will converge at the *upper* focus, marked as \mathbf{p} in the figure.

To derive the equations for hyperboloids that focus light at \mathbf{p} , we make use of two pieces of information. First, from classical geometric optics, we use the fact that for a material having index of refraction n , the hyperboloidal surface must have an asymptotic cone of half-angle $\alpha = \arccos(1/n)$. This allows us to express b in terms of n and a :

$$b = a \tan(\arccos(1/n)). \quad (6)$$

To make subsequent formulas more readable, we will write

$$t = b/a = \tan(\arccos(1/n)). \quad (7)$$

Next, we translate the hyperboloid such that its upper focus is at \mathbf{p} , making use of the fact that the distance from the center to a focus is $\sqrt{a^2 + b^2} = a\sqrt{1 + t^2}$:

$$\frac{(y + a\sqrt{1 + t^2} - p_y)^2}{a^2} - \frac{(x - p_x)^2 + (z - p_z)^2}{a^2 t^2} = 1. \quad (8)$$

As in the reflective case, once we are able to construct a family of hyperboloids we may select one that meets our requirements. We again impose the conditions of manufacturability and maximal shallowness, construct the best patch, taper its sides, and subdivide into segments.




Figure 5: Hyperboloid with \mathbf{p} as its farther focal point. The angle α is half the opening angle of the asymptotic cone.

3.2 Point Sampling

Because our generated highlight holograms will have a relatively small number of points, we focus on placing those points in meaningful locations: we wish to convey the surface shape as clearly as possible with an economy of visual markings. This problem is frequently encountered in the domain of *Non-Photorealistic Rendering*, specifically in the problem of (sparse) line rendering. We therefore build upon line-drawing techniques by constructing a probability distribution over the target mesh based on the likely locations of lines in sparse drawings, sampling from this distribution to yield an initial set of points, then selecting a subset of these that result in non-overlapping grooves. The initial sampling stage is described here, while the selection of grooves is described in the following subsection.

Over the past several years, researchers have developed a variety of mathematical definitions for placing lines on a 3D shape, typically using local differential properties of the surface [Saito and Takahashi 1990; Hertzmann and Zorin 2000; DeCarlo et al. 2003; Ohtake et al. 2004; Judd et al. 2007; Lee et al. 2007; DeCarlo and Rusinkiewicz 2007]. These methods produce aesthetically appealing drawings, in many cases match artist-drawn lines [Cole et al. 2008], and can be shown to convey 3D shape effectively [Cole et al. 2009]. Because different lines are appropriate for different models, and fully automatic algorithms cannot yet substitute for a user’s judgment, we leave the choice of algorithm(s) and appropriate parameters/thresholds to the user. For a majority of the results presented in the paper, we used occluding contours (interior and exterior silhouettes) and suggestive contours [DeCarlo et al. 2003], though the heptoroid example also includes principal high-light lines [DeCarlo and Rusinkiewicz 2007].

We begin by computing the locations of lines on the surface, for a dense sampling of viewpoints around the central one. Because the scratch holograms are designed for horizontal parallax, we sample viewpoints in a 30-degree horizontal range, and only a 5-degree vertical range. For each vertex in the model, we compute the fraction of viewpoints for which our line-drawing definitions placed a line within a face touching that vertex. Normalizing the per-vertex values over the mesh results in a probability distribution. (We assume that our meshes have close to uniform vertex spacing — we could extend our method to meshes with greater variance in face size by re-tessellating the mesh.) We also give the user the option of adding a small constant probability to each vertex, which results in additional points uniformly sampled over the mesh.

Once we have a probability distribution, we use stratified sampling over projected image coordinate (in the central viewpoint) to draw




Figure 6: Point sampling and packing pipeline. Starting from a 3D model (top), we compute a probability distribution (center), and select samples (below). Points resulting in non-overlapping grooves (white) are selected, while overlapping points (green) are possibly moved to a nearby point on the mesh (yellow). For some points (blue) no non-overlapping point in the nearby neighbourhood can be found, and completely-invisible points (red) are pruned directly.

samples from it. We typically generate thousands of samples, since the subsequent “packing” stage will eliminate many of them.

3.3 Point Packing

Once we have obtained a set of sample locations on the mesh, we convert them to grooves on the hologram surface. Using ray tracing, we determine the extent of viewpoints from which the sample should be visible, then compute the groove bounding rectangles and paraboloidal/hyperboloidal patches using the formulas derived in Section 3.1. (Any completely-invisible points are pruned.)

We must now select a non-overlapping subset of grooves from the candidates. To do this, we traverse the samples in random order, adding each groove to the output set if its bounding rectangle does not overlap any previously-added ones. If there is an overlap, we could simply drop the point, but to maximize the number of visible points we instead attempt to move to a nearby point on the mesh. To do this, we collect all vertices on the mesh within some distance k of our initial point \mathbf{p} :

$$N_k(\mathbf{p}) = \{\mathbf{q} \mid d(\mathbf{p}, \mathbf{q}) \leq k\}. \quad (9)$$

We assign a weight to each point based on its previously-computed probability and Euclidean distance to \mathbf{p} , sort the points in decreasing order, and traverse the list in order to find a non-overlapping replacement for the initial sample. Figure 6 illustrates this process.

3.4 Color

To obtain a color hologram from a model with per-vertex color, we look up the color at each point that resulted in a groove. These are used to construct a map consisting of colored rectangles at each groove location, as shown in Figure 7. This map is printed on a

transparency using an inkjet printer, then overlaid on the hologram (whether reflective or transmissive). We have also experimented with including shading (such as Lambertian reflection) in the computed color, which gives an even more compelling 3D illusion.




Figure 7: Part of the colormap from the santa hologram (left). Additional shading (right).

3.5 2D Stippled Images



Figure 8: Part of the image used to create the stippling (left). Point set (center). Image of the stippling facricated with a 3D printer (right).

We now consider a number of extensions of the method of milling curved grooves in a surface, in which the ultimate goal is not necessarily to convey a 3D shape. We begin with the simple goal of conveying a 2D image. While this is less complex than 3D “holography,” the technology is nevertheless compelling because it allows embedding a 2D image in an arbitrary surface (including non-planar ones), in a way that is compatible with modern methods of mass production (especially injection molding). In contrast, technologies such as printing and application of stickers have difficulties with curved surfaces, and frequently must include a manual component.

We begin by constructing a point set from a gray-scale image using the Weighted Voronoi stippling algorithm of Secord [2002]. Since all of these points are to lie at the surface, any highly-curved and

small groove is sufficient to create a highlight for most views and most directions of incident light. In practice, we use small hemispherical groove for each stippled point; all hemispheres have the same radius.

The same method can be used for both reflective and transmissive surfaces. Figure 8 shows a result for the reflective case. Although a closeup reveals a small amount of stray light in the fabricated result (at right), the use of short and shallow grooves in general results in high-quality and easily visible results. We therefore believe that this fabrication mode is production-ready, in contrast with our other results, which may still be considered experimental.

3.6 Animation

Our final use of fabricated grooves is to create 2D animations: the goal is to reveal different images to a stationary observer as the light source direction is changed. As input we use a set of animated meshes with vertex correspondence. For example, Figure 9 shows an input point set obtained by projecting 13 frames of a captured human animation [Vlasic et al. 2008]. For each mesh we create a point set by using the method described in section 3.2. The gathered vertex indices are merged and traversed in a random order to generate the grooves as follows.

For each vertex, we construct a potential groove for every frame in which the point is visible. We resolve overlappings between these grooves by merging them. This is done by deleting two overlapping (consecutive) grooves and constructing a new one with an adjusted field of view, such that the groove produces a visible highlight for both frames. We repeat this until no two grooves overlap. Then we add all grooves that do not overlap with the already added ones (which have another vertex index). To get a smoother transition between two frames, we add as many grooves as possible between two consecutive frames (using linear interpolation).

4 Results

In this section, we first describe the manufacturing process we used to produce our prototypes. Then, we present results for the entire pipeline, from 3D model to highlight hologram, for several different examples. Our results encompass reflective and refractive (transmissive) holograms, as well as ones that include color and shading. Finally, we describe limitations of the current process and propose ideas for removing them.

4.1 Fabrication using Engraving Machines

We use a Roland EGX-600 computer controlled engraver, shown in Figure 10, to fabricate our holograms. This machine offers high precision (0.01 mm), large working area (610 mm x 407 mm), high




Figure 9: Input point set for three of the thirteen frames of an animated sequence.



Figure 10: We used a Roland EGX-600 engraver for our experiments. This is a 3-degree-of-freedom milling machine with high-precision axes and a fine-tipped cutter (as shown in closeup at left).

manufacturing speed, ease of operation, and excellent bundled software (Modela Player and Virtual Modela). Our pipeline outputs a hologram surface as an STL file, and we use Modela Player to convert this surface to output used by the EGX-600. We also have an option to simulate the engraved surface using Virtual Modela. We use a high speed steel engraving cutter (model HS-171-005K) with a 0.005-inch tip, as shown in Figure 10, left.

We use acrylic plates (300 mm × 300 mm × 6.25 mm size) as the substrate for our highlight holograms. Acrylic is soft, easy to engrave, and can be easily polished to make the surface highly transmissive or specular. We set the machine resolution at 0.05 mm, and we generate holograms with an overall size of at least 250 mm. The groove width is set to 2 mm, allowing us to fit up to a thousand points on this surface. The average time to mill a surface of this size at this resolution is about 18 hours. In order to make the surfaces highly specular we paint the reflective surfaces using an air-brush with a black high gloss paint. We have also used black chrome plating on some of the results and we have obtained even better results. For the transmission hologram we have used a high gloss transparent lacquer.

4.2 Example Computational Highlight Holograms

We have generated computational highlight holograms from a number of popular 3D models. We produced a total of three reflective holograms, one transmissive, where one reflective hologram is with color. Further we fabricated two reflective stiplings, one transmissive, and one animated 2D image. Our results are shown in Figures 1, 13, and 14, while some statistics are summarized in Table 1.

Reflective holograms. We have fabricated three reflective hologram: the Grid, the Armadillo. The sizes of the hologram area on the 300 x 300 mm plate are: 295 x 145 mm, and 275 x 300 mm. The ideal viewer position is perpendicular to the plate area, with a field of view of 30 degrees in the horizontal direction, while the assumed light direction is 25 degrees above the viewer.

The groove length ranges from 2 mm to 36 mm (the longest grooves are for the 3D points that are very far behind the hologram surface and the shortest ones correspond to points that appear on the hologram surface). The groove width is relatively constant, ranging from 0.6 mm to 0.8 mm, while groove depth ranges from 0.4 mm to 0.8 mm, with most of the grooves less than 1 mm deep. Given the current groove width and the size of the surface plates, our algorithms had difficulty packing more than about 3,000 points on the hologram. For the Grid hologram, for example, we fit only 518 on the surface. For the Armadillo model we used 2173 out of 32538 original points.

Our holograms are clearly visible both in sunlight conditions and indoors using a point light source (please see results in the accompanying video). A sampling of observers could easily see the

Table 1: Parameters and statistics of five different computational highlight holograms. All dimensions are in mm.

Model	Type	Plate Size	Light Direction	FOV	Groove Length min/max/avg	Groove Width min/max/avg	Groove Depth min/max/avg	Virtual Depth min/max/avg	Packed Points	Original Points
Grid	Refl.	295 x 145	25°	30°	6.0 / 81.7 / 36.1	0.6 / 0.9 / 0.6	0.4 / 0.9 / 0.8	2 / 148 / 75	-	518
Santa	Refl.	295 x 285	25°	30°	2.5 / 77.0 / 15.5	0.6 / 1.1 / 0.7	0.2 / 0.9 / 0.7	1 / 148 / 55	2042	41543
Armadillo	Refl.	275 x 300	25°	30°	2.0 / 81.3 / 19.0	0.6 / 0.8 / 0.7	0.2 / 0.8 / 0.7	4 / 154 / 67	2173	32538
Max Planck	Trans.	210 x 290	10°	10°	6.0 / 43.0 / 13.5	1.4 / 1.9 / 1.7	0.7 / 4.3 / 1.6	2 / 180 / 80	807	1634
Tutanchamun	Refl. Stipp.	200 x 200	-	-	0.7	0.7	0.14	-	-	12000
Einstein	Refl. Stipp.	330 x 330	-	-	0.7	0.7	0.14	-	-	50000
Crane	Anim.	230 x 300	15 - 60°	5°	1.15 / 1.19 / 1.16	0.79 / 2.05 / 0.94	0.05 / 0.97 / 0.28	-	-	2542

models floating in front of or behind the surface, though we found that for most observers it was easier to perceive points behind the surface. In order to obtain better contrast between the points and the rest of the surface, we use an acrylic plate with a bronze tint.

Transmission hologram. We have fabricated and show one transmission hologram, the Max Planck bust. In this case, light illuminates the hologram from the back and passes through the surface, again resulting in highlights when viewed from specific angles. In order to maximize transmission, this hologram was manufactured in a clear acrylic plate. As a result, when placed in a window and viewed in sunlight, the 3D points appear extremely bright, and we had to reduce the camera exposure significantly when recording video footage. The facial features of the model are clearly visible, because the samples lie mostly along lines. Most of the parameters of this hologram are similar to those used for reflective holograms, though the grooves are, in general, deeper.

Color hologram. For the colored hologram we use the Santa model, and have experimented with color albedo alone, as well as color modulated by Lambertian shading. The hologram statistics are quite similar to the other reflective holograms, though we chose to insert additional points uniformly distributed on the surface, since these points help convey the color of large, featureless regions. Although we used a two layer technique (a transparency and a regular reflective hologram) to manufacture the color hologram, we are also currently investigating jetting the inks directly onto the hologram surface with a modified EPSON printer. We also believe that it would be straightforward to produce a colored transmission hologram.

We believe that the addition of color and shading information improves the ease of 3D perception in the holograms, since it adds additional cues about the shape. Informally, users also prefer colored holograms to white ones.

Stippled 2D images. We have produced two reflective stipplings using a 3D printer, the Mask of Tutanchamun, and the Einstein Portrait (shown in Figure 14). Their sizes are 200 x 200 mm and 330 x 330 mm. From the two grey-scale images we created point sets with sizes 12000 and 50000. The images induced by the fabricated surfaces are clearly visible from almost any viewing direction. Furthermore, they can be viewed under natural illumination as shown in the video.

Animated 2D images. We have also fabricated a reflective surface with “animated” 2D highlights corresponding to the Crane animation shown in Figure 9. The corresponding sequence depicts a significant motion of a person. The fabricated surface reproduces the animation sequence quite well when viewed under point light source illumination. A few photographs of the surface under different light direction are shown in Figure 14 and in the included video.

5 Discussion

In this section, we first discuss the effects of changing the illumination direction on virtual images. Second, we analyze the field of view of our fabricated holograms and compare it to the theoretical limit. Finally, we discuss the limitations of computational highlight holograms.

5.1 Distortion

When analyzing our holograms, we have observed that changing the light direction when the observer is fixed has a similar effect to moving the observer in horizontal direction while the illumination is fixed. However, unlike changing view, changing the light direction distorts the virtual object.

In order to calculate the trajectory of a virtual point as a function of light direction, we assume the observer is parallel to the hologram plane and that the distance between his eyes is fixed. Based on the direction of the light source and the shape of the groove, we can numerically determine the position of the highlight in the image plane for each eye. Then the virtual point position can be computed by intersecting the rays that connect each eye with the corresponding highlights. However, computing the analytical expression for the 3D trajectory of the virtual point as a function of light direction requires finding roots of quartic equations. Therefore, we have decided to analyze the distortion numerically using typical parameter values of our fabricated holograms. In particular, we set the distance between the eyes to 60 mm, the size of the image plane to 300×300 mm with the center at $(0, 0, 0)^T$ and the viewer at $(0, 0, 1000)^T$.

In Figure 11 we show a grid of points in the xy plane and the point trajectories induced by a discrete set of light directions. The range of light directions is ± 15 degrees off-axis. As is clearly observable, the points move on curved trajectories. The shape of these trajectories depends both on the depth of a point (its y coordinate) and the relative difference between the viewer position and the initial point position in the x direction. We observe that points below the image plane move on “hyperbolic” curves and points above the image plane move on “elliptical” curves. In the case of the grid shown in Figure 11, the depth of each point can change up to 15%.

We also show that changing the light direction causes distortions that do not preserve angles and do not preserve straight lines. In order to demonstrate this we use a hologram with a model of a rectangle when illuminated on-axis (Figure 12). When the hologram is illuminate off-axis (at an angle of 30 degrees), it can be clearly observed that straight lines do not remain straight, the length of the lines is changed, and angles are not preserved. For example, the maximum change in length is 53% and the maximum angle distortion is 26 degrees. The distortion along the z axis is very low compared to the distortions in the xy plane. Both in the simulation and when viewing the fabricated holograms this distortion has no significant effect.




Figure 11: A hologram with a uniform grid of points when illuminated on-axis. As the light direction departs from the initial setting (± 15 degrees), the points move on curved trajectories. The depth of each point can change up to 15%.




Figure 12: Distortions due to off-axis illumination introduce distortions that do not preserve angles, lengths, and straight lines. The initial rectangle has side length 250 mm and depth 200 mm (shown in black). When illuminated off-axis with an angle of 30 degrees, the rectangle is distorted (red).

5.2 Field of View

Vertical Field of View. Using the same notation as described in section 3.1 we can analyze the theoretical vertical field of view of a groove. The expression for this field of view is:

$$\alpha = 2 \arctan \left(\frac{\|\mathbf{v} - \mathbf{p}\| w}{2\|\mathbf{v} - \mathbf{i}\|\|\mathbf{p} - \mathbf{i}\|} \right), \quad (10)$$

where \mathbf{v} is the viewer position, \mathbf{p} the location of the 3D point, w is the groove width, and \mathbf{i} is the intersection of the line passing through \mathbf{v} and \mathbf{p} with the hologram plane. Whereas the theoretical vertical field of view is usually very small (about 1 degree), the observed vertical field of view of the fabricated holograms is about

10 degrees. This is due to highlight broadening resulting from a glossy, but not perfectly specular, milled surface.

5.3 Limitations

In experimenting with prototype computational highlight holograms, we have observed some of the limitations of the technique. First, it is difficult to reproduce horizontal lines due to the groove overlap (which we always forbid). Reducing groove size and jittering or offsetting the points slightly can alleviate some of this problem. Second, we have observed that the direct light reflection from the unmilled portion of the surface is often distracting for hologram viewing. We have found that roughening the unmilled area, making it more diffuse, helps to ameliorate this problem. We also design our holograms for angled incident light, helping ensure that the direct specular reflection is not visible for typical viewer positions. Third, the field of view of the hologram is limited in the vertical direction. Increasing the relative groove width would help with this problem, at the cost of not allowing as many grooves to fit on the surface. We also noticed that scenes might appear slightly distorted from the extreme viewpoints in the horizontal direction. We believe it is possible to resolve this problem by subdividing the grooves into vertical segments.

6 Conclusion

In this paper we have presented an investigation of computational highlight holography — an automated process that converts a three-dimensional computer model into surfaces that reproduce images of the model with correct binocular and motion parallax when illuminated by a directional light source. We show how to design, simulate, and manufacture reflective and transmissive surfaces and how to add color to these holograms. We believe that this method has a tremendous commercial potential due to its simplicity and flexibility.

There are many possible directions for future work. First, we would like to reduce the size of the geometric primitives in order to pack at least two orders of magnitude more points into the same surface. We believe that using around 100,000 points would allow us to represent the surfaces in a more continuous manner (as opposed to depicting only the most relevant surface features). Second, we plan to extend this method to curved surfaces and embed holograms into surfaces such as a car or a wine glass. This would increase the applicability of this method. We also would like to produce extremely large holographic surfaces. We believe that this is possible even using our current equipment, e.g., by dividing the surface into tiles. Finally, we plan to experiment with the shape of the surface fragments in order to produce grooves with different relative intensity. This, in turn, would allow us to produce holograms with appropriate shading without additional overlays or inks.

As a final note, we believe that there are many interesting techniques that are similar in spirit to computational highlight holography and that are stuck in the “analog” world. It would be beneficial to explore these methods and upgrade them to the computational era.

Acknowledgements

We thank Cyberware, Viewpoint, Stanford, and MPI for the 3D models used in this paper, and acknowledge NSF grant #CCF-0702580 for funding. The portrait of Einstein (Fig. 14) is retrieved from wikimedia.




Figure 13: Three highlight holograms produced using our process. For each hologram we show the model, and photographs of the fabricated hologram from the outer left view point (third row) and outer right view point (fifth row). The Max Planck hologram (leftmost) is transmissive, while the others are reflective.




Figure 14: Two 2D stippled images are shown in the first and the second row. The images on the left show the input image. The images on the right show the corresponding photographs of the fabricated surfaces under natural illumination conditions. The last row shows six photographs of the fabricated surface that represents an animation sequence.

References

- ABRAMSON, N. H. 2000. Incoherent Holography. *Proc. SPIE (Holography 2000)*, Vol. 4149, No. 1 (Oct.), 153–164.
- BEATY, W. J., 1995. Hand-Drawn Holograms. <http://amasci.com/amateur/holo1.html>.
- BEATY, W. J. 2003. Drawing Holograms by Hand. *Proc. SPIE (Practical Holography XVII and Holographic Materials IX)*, Vol. 5005, No. 1, 156–167.
- BENTON, S. A. 1969. Hologram Reconstruction with Extended Light Sources. *Journal of Optical Society of America*, Vol. 59 (Oct.), 1545–1546.
- COLE, F., GOLOVINSKIY, A., LIMPAECHER, A., BARROS, H. S., FINKELSTEIN, A., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2008. Where Do People Draw Lines? *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 27, No. 3 (Aug.), 88:1–88:11.
- COLE, F., SANIK, K., DECARLO, D., FINKELSTEIN, A., FUNKHOUSER, T., RUSINKIEWICZ, S., AND SINGH, M. 2009. How Well Do Line Drawings Depict Shape? *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 28, No. 3 (Aug.), 28:1–28:9.
- DECARLO, D., AND RUSINKIEWICZ, S. 2007. Highlight Lines for Conveying Shape. In *Proc. NPAR*, 63–70.
- DECARLO, D., FINKELSTEIN, A., RUSINKIEWICZ, S., AND SANTELLA, A. 2003. Suggestive Contours for Conveying Shape. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 22, No. 3, 848–855.
- DENISYUK, Y. 1962. Photographic Reconstruction of the Optical Properties of an Object in Its Own Scattered Radiation Field. *Sov. Phys. Docl*, Vol. 7, 543–545.
- EICHLER, J., DÜNKEL, L., AND GONÇALVES, O. 2003. Three-Dimensional Image Construction by Curved Surface Scratches. *Applied Optics*, Vol. 42, 5627–5633.
- FUCHS, M., RASKAR, R., SEIDEL, H.-P., AND LENSCH, H. P. A. 2008. Towards Passive 6D Reflectance Field Displays. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 27, No. 3 (Aug.), 58:1–58:8.
- GABOR, D. 1948. A New Microscopic Principle. *Nature*, Vol. 161 (May), 777–779.
- GARFIELD, E. 1981. ISI's 'World Brain' by Gabriel Liebermann: The World's First Holographic Engraving. *Essays of an Information Scientist*, Vol. 5, 348–354.
- HERTZMANN, A., AND ZORIN, D. 2000. Illustrating Smooth Surfaces. In *Proc. SIGGRAPH*, 517–526.
- IVES, F., 1903. Parallax Stereogram and Process for Making Same. U.S. Patent No. 725,567.
- IVES, F. 1928. A Camera for Making Parallax Panoramagrams. *Journal of the Optical Society of America*, Vol. 17 (Dec.), 435–439.
- JONES, A., MCDOWALL, I., YAMADA, H., BOLAS, M., AND DEBEVEC, P. 2007. Rendering for an Interactive 360° Light Field Display. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 26, No. 3 (July), 40:1–40:10.
- JUDD, T., DURAND, F., AND ADELSON, E. 2007. Apparent Ridges for Line Drawing. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 26, No. 3, 19:1–19:8.
- LEE, Y., MARKOSIAN, L., LEE, S., AND HUGHES, J. F. 2007. Line Drawings via Abstracted Shading. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 26, No. 3, 18:1–18:8.
- LEITH, E., AND UPATNIEKS, J. 1962. Reconstructed Wavefronts and Communication Theory. *Journal of the Optical Society of America*, Vol. 52, No. 10 (Oct.), 1123–1130.
- LIPPMANN, G. 1908. Epreuves Reversibles Donnant la Sensation du Relief. *Journal of Physics*, Vol. 7, No. 4 (Nov.), 821–825.
- MATUSIK, W., AND PFISTER, H. 2004. 3D TV: A Scalable System for Real-Time Acquisition, Transmission, and Autostereoscopic Display of Dynamic Scenes. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 23, No. 3 (Aug.), 814–824.
- OHTAKE, Y., BELYAEV, A., AND SEIDEL, H.-P. 2004. Ridge-Valley Lines on Meshes via Implicit Surface Fitting. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 23, No. 3 (Aug.), 609–612.
- OKOSHI, T. 1976. *Three-Dimensional Imaging Techniques*. Academic Press, New York.
- PATOW, G., AND PUEYO, X. 2005. A Survey of Inverse Surface Design from Light Transport Behavior Specification. *Computer Graphics Forum*, Vol. 24, No. 4 (Dec.), 773–790.
- PATOW, G., PUEYO, X., AND VINACUA, A. 2007. User-Guided Inverse Reflector Design. *Computers & Graphics*, Vol. 31, No. 3 (June), 501–515.
- PERLIN, K., PAXIA, S., AND KOLLIN, J. S. 2000. An Autostereoscopic Display. In *Proc. SIGGRAPH*, 319–326.
- PLUMMER, W. T., AND GARDNER, L. R. 1992. A Mechanically Generated Hologram. *Applied Optics*, Vol. 31, 6585–6588.
- SAITO, T., AND TAKAHASHI, T. 1990. Comprehensible Rendering of 3-D Shapes. In *Proc. SIGGRAPH*, 197–206.
- SECORD, A. 2002. Weighted Voronoi Stippling. In *Proc. NPAR*, 37–43.
- VLASIC, D., BARAN, I., MATUSIK, W., AND POPOVIĆ, J. 2008. Articulated Mesh Animation from Multi-view Silhouettes. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 27, No. 3 (Aug.), 97:1–97:10.
- WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating Microgeometry for Custom Surface Reflectance. *ACM Trans. Graph. (Proc. SIGGRAPH)*, Vol. 28, No. 3 (July), 32:1–32:6.