

Dispersion-based Color Projection using Masked Prisms

Rafael Hostettler¹ Ralf Habel¹ Markus Gross¹ Wojciech Jarosz^{2,1}

¹Disney Research Zurich ²Dartmouth College




Figure 1: Our approach is able to project a full color image (Input) using a device which contains no colored components—consisting only of two masks (2nd, hidden), and one or two prisms. (The central panel shows the effect of displaying the Input image on an LCD display for comparison)

Abstract

We present a method for projecting arbitrary color images using a white light source and an optical device with no colored components—consisting solely of one or two prisms and two transparent masks. When illuminated, the first mask creates structured white light that is then dispersed in the prism and attenuated by the second mask to create the color projection. We derive analytical expressions for the mask parameters from the physical components and validate our approach both in simulation and also demonstrate it on a wide variety of images using two different physical setups (one consisting of two inexpensive triangular prisms, and the other using a single rhombic prism). Furthermore, we show that optimizing the masks simultaneously enables obfuscating the image content, and provides a tradeoff between increased light throughput (by up to a factor of three) and maximum color saturation.

1. Introduction

Recently, many methods and devices have been proposed that use optical elements or optical material properties in unusual and intriguing ways to reveal hidden images or create unexpected and surprising effects through shading/shadowing [AM10,BBAM12,MP09], refraction/reflection [FDL10,PJJ*11], or attenuation [WLHR11,BKB*12] to name a few examples. In this spirit, we create a device that projects full-color images out of white light without having any colored part, consisting only of two greyscale masks attached to one or two prisms. The source of the color thus remains initially unclear until the viewer realizes that white light consists of different wavelength components that are decomposed by the prism(s) through dispersion and filtered by the masks. Like previous work, our device provides intriguing applications in art and entertainment, as well as

education—by providing an engaging way to demonstrate and teach the physics of dispersion to children.

The input to our method is a full-color image (Figure 1 far left), which we wish to produce using our optical device. We present a simple computational approach to compute two mask patterns (Figure 1 left & center) which, when printed on transparencies and placed on opposite ends of our prism(s), lead to color image projections due to dispersed white light (Figure 1 right).

At a high-level our setup shares some similarities with traditional LCD displays/projector which contain red, green and blue sub-pixels, but we attain this using a different optical process. We also conceptually create sub-pixels using our first mask; however, instead of obtaining sub-pixels with simple color filters as in LCDs, we induce a repetition of small “rainbow” sub-pixels when light passes through our

first mask and disperses through the prism(s). These rainbows are then attenuated by the second mask to reproduce weighted combinations of the constituent colors for each pixel.

We derive a simple analytic model to compute the two masks to reproduce arbitrary color images. Furthermore, we show that by numerically optimizing the masks, we can obtain higher light throughput by sacrificing maximum color saturation. The same optimization can be used to obfuscate the image content so that the source image is harder to infer from the masks themselves.

We validate our method in simulation and demonstrate it using two physical prototypes. Our approach is practical, easy to reproduce, and inexpensive (our multi-prism prototype uses standard 60° triangular prisms and costs under \$30).

2. Related Work

Our work can be categorized as a method in the growing field of computational materials and displays, where basic principles of optics and light transport are used and optimized in novel ways. In contrast to this previous work, our approach relies on the concept of dispersion in prisms.

Computational Materials & Displays. An early example of computational materials is the generation of bas-reliefs from 3D scenes [WDB*07]. Alexa and Matusik [AM10] fabricated diffuse surfaces which are capable of producing two images when illuminated from different directions. This was further extended by Bermano et al. [BBAM12] to produce up to four distinct images. Mitra and Pauly [MP09] created 3D objects that cast prescribed shadows for different incident light directions.

Using refraction, Fuchs et al. [FRSL08] created a light-dependent reflectance field display using lenslet arrays. Weyrich et al. [WPMR09] fabricated surfaces with custom micro-facet distributions capable of producing simple stippled images as reflective caustics. More recent techniques have fabricated custom lenslet arrays [PJJ*11] or rearrangeable sticks [YIC*12] to produce monochromatic images as refractive caustics. Similarly, Finckh et al. [FDL10] optimized reflectors and refractors to produce desired caustic images in simulation. Many of these approaches rely on idealized point or directional light sources, similar to our approach.

Wetzstein et al. [WLHR11] used multi-layer attenuators, together with a uniform backlight, to create light-field displays of static objects. Baran et al. [BKB*12] applied this concept in reverse to cast distinct color shadows when illuminated from different directions.

Naor et al. [NS95] use transparent masks to encrypt an image which can only be recovered if the masks are exactly positioned over each other. Our method also uses two masks to represent an image, which we can obfuscate using random initialization, though we do not currently consider cryptographic strength.

Prisms and Multi-spectral Projection. Prisms are one of the most basic devices in optics and spectroscopy, and are one of the main component for analyzing and modifying light paths since the beginning of optics. Prisms can also be used in sequence to amplify their dispersive effect or re-collimate the dispersed light, allowing for parallel light output at the last refraction (which we exploit in our setup). Prisms have also been used as beam expanders in ultrafast tunable lasers [Dua03]. We also use prisms as a form of beam expander but to induce a series of modulated rainbows to form our projected images.

Cao et al. [CDT*11] recently used the standard slit-prism spectrometer setup in an array fashion by applying a binary mask to a prism for low-resolution multi-spectral video acquisition. As the light field passes through the slits and prism it gets dispersed into its spectral components, which are then captured by a monochromatic camera. Our approach shares similarities to this method, but operates in reverse to project color images using several masks.

Mohan et al. [MRT08] use a diffraction grating to split the light of a projected *color* image into its spectral components, which are then modulated individually, before being re-integrated into a modified color image. Though not shown in their work, a prism could also act as a dispersive element instead of the diffraction grating. We do not filter the spectrum of a colored image but directly create an image from white light.

Lin et al. [LWLD14] use a snapshot approach to hyperspectral imaging where two high-speed spatial light modulators, located conjugate to the image and spectral plane, respectively, code the hyperspectral datacube into a single sensor image such that the high-resolution signal can be recovered in postprocessing.

Besides fulfilling a different purpose than our approach, these methods rely on complex optical setups with many expensive components that need to be calibrated, whereas our approach only consists of one or two prisms with simple masks attached. Furthermore, diffraction gratings induce significantly more light loss (greater than 30% compared to less than 2% in a prism) and always divert light into two directions instead of just one as is the case with a prism.

3. Method Overview

Our goal is to project color images using white light, two masks, and one or two prisms. We illustrate the principal configuration of the one-prism and two-prism setups in Figure 2 and show the physical prisms with masks attached in Figure 3. Our experimental setup (Figure 4) consists of a directional light source, the masked prism, and a projection surface.

To obtain our masks, we first mathematically formulate how an incident light path is dispersed into a rainbow by our prism setup (Section 4.1). We use this information to then design our first mask pattern (Section 4.2) to induce a good




Figure 2: The light paths considering dispersion in a rhombic prism (top), and the triangular two-prism configuration (bottom)

set of rainbow basis functions. These basis functions are then weighted by our second mask (Section 4.3) to obtain our final image.

Since we are not limited to fix the set of rainbow basis functions created by the first mask, we can optimize both mask to produce a final image. We can optimize a set of masks for a higher light throughput (gaining up to a factor of 3) by introducing limits on the maximum color saturation. The same optimization can be used to obfuscate the image content in the masks so that the image content is difficult to infer from the masks alone.

We show a variety of results in Section 5 and discuss the benefits and limitations of our approach as well as ideas for future work in Section 6.

4. Approach

We consider two geometric configurations: a two-prism and single-prism setup. In both setups our prisms are made of PMMA, whose wavelength-dependent index of refraction




Figure 3: The rhombic prism (left) and the two triangular equilateral with masks attached.




Figure 4: Experimental setup consisting of a directional light source (left), the masked prism (middle right) and the projection area (right). Blockers avoid stray light near the projection area.




Figure 5: The wavelength-dependent index of refraction, $\eta_i(\lambda)$, for Polymethyl Methacrylate (PMMA).

$\eta_i(\lambda)$ can be found in standard optics references [KSIN07]. We plot $\eta_i(\lambda)$ for PMMA in Figure 5.

Our two-prism setup (see Figure 6) uses two standard equilateral triangular prisms (interior angle $\beta = 60^\circ$) where the two inner as well as outer faces along the light path are parallel. The outer faces of the prisms are separated by a distance d_0 . The incident directional illumination arrives at the first prism with an incident angle δ_0 . In this setup, light passes through 4 interfaces (with five corresponding angles $\delta_0 \dots \delta_4$) before exiting the second prism with direction δ_4 .

Our one-prism setup uses a single rhombic (parallelepiped) prism with an inner angle of 60° (see Figure 2). This can be interpreted as two fused equilateral triangular prisms without the two interior refractions. We therefore describe and derive the two equilateral prism setup in detail and the rhombic prism follows as a simplification.

4.1. Light Propagation

The path of light traveling through a refractive boundary is described by Snell's Law:

$$\eta_i(\lambda) \sin(\delta_i) = \sin(\delta_o), \quad (1)$$

where δ_i and δ_o are the angle of light inside and outside the material, and η_i is the relative wavelength-dependent index of refraction of the two materials.




Figure 6: Two equilateral prism setup. The light enters the first prism at an angle δ_0 , gets refracted 4 times ($\delta_{1..4}$) and exits with the same angle it enters.

Given the geometric configuration of our prism(s) ($d_0, d_1, \beta, \eta_i(\lambda)$), we wish to determine the aggregate effect of this optical setup on light incident on the first prism at an angle δ_0 . In particular, we wish to determine the final outgoing direction δ_4 , as well as the spatial offset o_y along the y -axis (see Figure 6).

Angular Domain. Due to the parallel planes induced by our prism setup, it turns out that $\delta_4 = \delta_0$. In effect, for any light path that makes it through all four interfaces the optical device does not change the light ray direction, for any wavelength. This is easy to show by repeated application of Snell's Law to the two sets of parallel planes in Figure 6. Less formally, each set of parallel interfaces effectively cancels out the angular effect of refraction.

Spatial Domain. Refraction through parallel planes does however induce a spatial shift of the outgoing ray. Furthermore, the light ray offset is *dependent* on the material's index of refraction because different wavelengths take different paths through the setup due to *dispersion*. If we just consider a single point of incoming white light entering the setup (see Figure 6), its spectral content gets dispersed over a small line segment, creating a tiny rainbow on exiting the second prism. Though complicated, the offset can be calculated analytically by applying geometric optics. We set the outer parallel surfaces as the principal y -axis as the coordinate system for the calculation (see Figure 6). By applying Equation (1) consecutively, the offset o_y of the exit point is:

$$o_y(\lambda) = \frac{d_1 \cos \left(\beta \arcsin \left(\eta_i(\lambda) \cos \left(\beta + \arcsin \left(\frac{\sin \delta_0}{\eta_i(\lambda)} \right) \right) \right) \right) \sin \beta}{\sqrt{1 - \eta_i^2(\lambda) \cos^2 \left(\beta + \arcsin \left(\frac{\sin \delta_0}{\eta_i(\lambda)} \right) \right)^2}} + \frac{(d_0 - d_1) \sin \delta_0}{\eta_i(\lambda) \sqrt{1 - \frac{\sin^2 \delta_0}{\eta_i^2(\lambda)}}} \quad (2)$$

Note that the offset is *independent* of the actual incident position of light; hence, light incident at any position on the incident surface is offset by the same amount as long as it crosses all 4 interfaces.

Dispersion Extent. Equipped with this measure, we can directly calculate the extent e of the dispersion by simply taking the difference of the minimum and maximum considered wavelengths:

$$e = o_y(\lambda_{min}) - o_y(\lambda_{max}). \quad (3)$$

We consider the typical bounds of the human visual system and hence set $[\lambda_{min}, \lambda_{max}] = [390 \text{ nm}, 750 \text{ nm}]$. For PMMA, this results in $\eta_i(\lambda_{min}) \approx 1.53108$ and $\eta_i(\lambda_{max}) \approx 1.485139$.

Determining the Setup. The free parameters for our setup include the incoming light directions δ_0 , and, in the case of the two-prism configuration, d_1 (is then strictly determined from the size of the prisms).

We set δ_0 so that the refracted light path for a wavelength of 550 nm (the center of our considered range) is parallel to the adjacent prism side, giving us a $\delta_0 = 48.3^\circ$ for the index of refraction for PMMA. Though our method is insensitive to the exact value of δ_0 , this strategy maximizes the size of the projection due to the finite extent of the prisms.

In the case of the two prism setup, we separate the prisms by a distance d_1 . Varying d_1 influences e , and therefore changes the resolution of the projection as we will discuss in the following section. As mentioned previously, the rhombic prism setup corresponds to setting $d_1 = 0$.

4.2. Defining the First Mask

If we consider directional white light incident across the entire surface of the first prism (without any masking), we obtain a white image as the projection since the prism(s) only slightly shift the spectral content by $o_y(\lambda)$.

Our goal is to create pixels in the projection, each consisting of colored sub-pixels like in an LCD projector, but induced by dispersing white light into its spectral components. Hence, each pixel (or basis function) in our projector will be a small rainbow.

The most naive way to induce a set of pixels with rainbow sub-pixels is to use an impulse train as the first mask. As light passes through our optical setup, this impulse train gets convolved by the offset function $o_y(\lambda)$ of the prism to create a repeating pattern of rainbows. To get tightly-packed repeating rainbows, the spacing between the impulses needs to be e .

This naive approach has the benefit of representing spectral colors very accurately. The problem, however, is that infinitesimal impulses allow only an infinitesimal amount of light to pass through, resulting in poor light efficiency. Our approach, however, allows a continuous trade-off between intensity and

spectral resolution: by increasing the size of the impulses to finite boxes, we let more light through (increasing light efficiency), but the resulting rainbow spectra are convolved with a box filter (lowering spectral resolution).

In our prototypes, we want to display 3-color images, so we choose a slit size of $e/2$ with a gap of size e between the slits (see Figure 7 top). This effectively convolves the spectrum of the light source with a box function of size $e/2$, smoothing out possible spikes in the spectrum. By convolving the incident spectral lightfield with this slit pattern, the resulting pixels have a width of $3/2e$, with a continuous (box-filtered) spectrum as the sub-pixels. A close-up visualization of the resulting idealized basis functions (simulated using spectral ray-tracing) is shown in Figure 7 bottom.




Figure 7: A part of the first mask (top) and a simulated projection (bottom) using spectral ray-tracing.

4.3. Defining the Second Mask

In order to create a color image projection, the second mask (placed on the light-exiting surface) needs to mask out different parts of the spectra to create the desired pixel color values. In theory, we have an infinite resolution along the image axis which is orthogonal to the dispersion direction; however, since most images have a pixel aspect ratio of 1, we chose the same resolution for both axes to form square pixels. We divide each pixel into 3 sub-pixel areas corresponding to 3 spectral box functions covering the red, green and blue areas of the spectrum. Our second mask modulates each corresponding sub-pixel with the linearized red, green and blue value from our input image.

The second mask continuously weighs the color contributions of the rainbow basis functions. We however produce our physical masks using a Laser printer, which approximates the continuous greyscale tones using binary dithering. To incorporate the effect of printer dithering, we show stochastically dithered masks in our simulated results.

To assess the expected fidelity of the pixels, Figure 8 shows a close-up (with projection area $\approx 8.4 \times 8.4$ mm) of a small test image of saturated colors, as well as the stochastically dithered mask, and the resulting simulated and physical result.

4.4. Optimizing Masks

The previously described approach induces rainbow pixels with the first mask, allowing to combine it with any second




Figure 8: A 9×9 pixel input image (left), the stochastically dithered mask (middle left), simulated image (middle right) and a macro photograph of a physical projection of size 8.4×8.4 mm (right).

mask that encodes an image. Instead of choosing a train of rainbow pixels as the initial basis functions for each row of an image, we can optimize both masks to produce an image and optimize for different goals, such as intensity throughput or obfuscation of the image content in the masks.

The train of rainbow pixels can be interpreted as a set of orthogonal basis functions, they do not overlap and therefore a final pixel color is only dependent on one rainbow (see Figure 7). In fact, the gap defined by e masks out 2 additional slits (of the width $e/2$) to the right of the one creating the rainbow. If we remove the gap, dividing it up into 2 additional slits, we triple the number of rainbow basis functions by overlapping them (see Figure 9). Rather than setting each triple of sub-pixels to (1,0,0) in the first mask (as done with the train of rainbow pixels), we now have 3 sub-pixel weights that can be arbitrarily chosen. Since rainbow basis functions overlap, this forms an overcomplete basis set and means that the red, green and blue contributions for a pixel no longer come from one, but from up to three sub-pixels in the first mask. The second mask still weights the different sub-pixels like in the unoptimized case. Since dispersion occurs




Figure 9: A part of an optimized set of masks where each sub-pixel a_m in the first mask induces weighted overlapping rainbows which are then weighted by the second mask sub-pixels b_m .

only in the horizontal direction, we can treat each scan line independently.

To formalize this, we denote a pixel P_n , with $n = 1 \dots N$ and where N the total number of pixels in a scan line (Figure 9). We denote the first mask A and the second mask B . We enumerate sub-pixels in the first mask as a_m with $m = 1 \dots 3N$ and similarly as b_m for sub-pixels in the second mask. The sum of a color channel in the triple of sub-pixels (forming one color pixel) should then match the color channel of the input image.

If we consider pixel P_n , instead of having a dedicated red sub-pixel, all 3 sub-pixels (a_m, a_{m+1}, a_{m+2}) of P_n in mask A can contribute to the red color of the pixel. In the case of green, contributions come from sub-pixels a_{m-1}, a_m, a_{m+1} , where a_{m+2} already contributes to the green component of the adjacent pixel P_{n+1} . For blue, the contributions are from a_{m-2}, a_{m-1} and a_m . The final color of a pixel, $P_{n(r,g,b)}$, after the light passes through both masks is therefore:

$$P_{n,r} = a_m b_m + a_{m+1} b_{m+1} + a_{m+2} b_{m+2} \quad (4)$$

$$P_{n,g} = a_{m-1} b_m + a_m b_{m+1} + a_{m+1} b_{m+2} \quad (5)$$

$$P_{n,b} = a_{m-2} b_m + a_{m-1} b_{m+1} + a_m b_{m+2} \quad (6)$$

with $m = 3n$.

For a given input image row $I_{n(r,g,b)}$, we want to find the weights of all sub-pixels in the two masks A and B to minimize:

$$\underset{A,B}{\operatorname{argmin}} \sum_n (I_{n,r} - P_{n,r})^2 + (I_{n,g} - P_{n,g})^2 + (I_{n,b} - P_{n,b})^2 \quad (7)$$

in order to project the final image with the masks. Equation (7) is a highly underdetermined problem and therefore many global and local minima exist. In fact, the train of rainbows is a global minimum of the stated problem.

We minimize Equation (7) using Levenberg-Marquardt optimization, minimizing over the vector of pixel color squared differences. Though a global minimum is not guaranteed, in practice we found that the large number of global and local minima assures a result very close to the optimal solution.

With this more general approach, we gain the possibility to optimize for other goals than simply the input image. In order to maximize the light throughput for example, we can optimize for an input image that contains larger values than 1[†], effectively increasing the overall intensity of the projected image. As a consequence, the range of displayable colors shrinks since the maximum intensity of fully saturated colors is 1. Yet, the minimization finds the best color reproduction for a given image. In general, a factor of 1.5–2 can still display an image faithfully, while the maximum possible value of

3 already strongly diminishes the quality of the projected image.

However, this highlights a fundamental difference between our approach and other commonly used ways to create projected color images: RGB filters reduce the light intensity a-priori by 2/3 as each of the three filters has to filter out the wavelengths accepted by the other two filters. This is the case for 3CCD projectors as well as color-wheel based DLP projectors. This is not the case for our approach, allowing for intensities higher than 1.

In order to obfuscate the image content, we want to find one of the other minima of (7) where the image structure is less apparent in the masks themselves. We can achieve this by choosing random initial conditions for a_m, b_m . Due to the large number of minima, the optimization will converge to a local minimum that retains the properties of the initial conditions, resulting in masks that are essentially noise.

5. Evaluation & Results

In this section we evaluate our approach in simulation, and demonstrate physical results.

5.1. Fabrication & Calibration

For our two-prism prototype, we use triangular prisms with dimensions $4 \times 4 \times 10$ cm which can be purchased in science and novelty shops. For our single-prism prototype, we use a custom-produced prism measuring $12.5 \times 9.2 \times 10.5$ cm with an inner angle of 60° . All our prisms are made of PMMA.

We print our masks using a Brother MFC-7320 Laser printer, and rely on the printer's internal dithering algorithm (we show simulated binary-dithered images for all our second masks, and continuous versions in the supplemental for direct printing).

After printing the masks, we need to attach them to the prism surfaces. Since we print sub-millimeter structures onto the masks, they need to adhere flat to the surface of the prisms. We use thin glass plates the size of the prism faces with a thickness of 2 mm to mount the masks. We secure the masks between the glass and prism by simply using rubber bands around the prism (see Figure 3). This approach has the advantage that it is easy to move the masks around during calibration, after which point they remain fixed. The additional thin glass plates introduce negligible dispersion and do not change the directionality of the incident light since each plate constitutes two parallel interfaces. On the large prism, the filters stick through static attraction, not requiring any additional mean of fixation.

We also conducted experiments with masks normally used for photo-lithography. These films are thicker, thus not requiring the extra glass for fixation. Here we cannot rely on printer dithering as these masks are produced binary. There

[†] 1 being the maximally achievable intensity for standard color projection methods, where white is created from three subpixels, filtering 2/3 of the light by construction.

was no visible improvement, hinting that other effects such as misalignments and the imperfect light source play a greater role than the dithering process, filter material and extra glass plate.

We place the setup so that it forms the angle δ_0 with the directional light source. Then, the masks are aligned by moving them slightly between the prism and the thin glass plate. Since the first mask consists of slits with the same size, we can arbitrarily move the second mask along the slit direction and we only need a relative alignment in the orthogonal direction, losing a pixel or two at the border due to the unknown slit offset. Though our setup construction is extremely simple, we nonetheless found that it allows for accurate alignment and once the positions are found, the glass plates can be fixed with tape to avoid movement of the masks.

5.2. Non-ideal Light Sources

In our derivation, we used an ideal directional light source, though many light sources have an angular spread which smooths out the resulting projection. Fortunately, the smoothing is not strong for small angles. Figure 10 shows the smoothing with an angular spread of 0.5° , which corresponds to the solid angle of the sun.

Also, the proposed method is dependent on the spectrum of the incoming light, but since the light spectrum gets convolved with a large box function (see Section 4.2), even fluorescent light sources which contain mostly spikes pose no problem for our approach. For the same reason, the color reproduction is only slightly dependent on the spectrum of the light source. We therefore do not perform full color calibration to the different spectra of our light sources (though this could improve color reproduction slightly). The slight color shifts introduced by a non-constant spectrum in our current physical prototypes are far smaller than from other artifacts such as misalignments and blurring.

5.3. Simulated & Physical Results

We test our approach using various images (see Figures 1, 12, 13) to assess resolution, color reproduction, and overall quality of the projection for both real and simulated results. We also discuss sources of errors for the case of simple and




Figure 10: Simulated results in the ideal case (top) and using an extended light source with an angular spread of 0.5° (bottom).

optimized masks. Our simulated results appear darker than the input since we visualize the individual R, G, and B components. Averaging each R, G, and B band into a single RGB pixel would result in images corresponding to the color input image.

To produce our physically-captured results, we illuminate our prism setup using white light from a standard Xenon lamp beammer placed at a distance of 4 m, producing nearly parallel incident light. We also succeeded using an LED flash light and sunlight. The principal size and resolution of the projections is directly given by the size of the incident surface area and the pixel size. With a pixel size of about 1 mm^2 , this results in a $40 \times 70 \text{ mm}$ projection at a resolution of 40×70 pixels in the double prism setup and a projection of $10 \times 10 \text{ cm}$ at a resolution of 100×100 pixels for the single-prism setup in the ideal case.

As the beammer is not a point-light source, a slight blurring and a consequent loss of contrast is introduced (see previous section and 10). Furthermore, as the light is also not perfectly parallel, the second filter has to be slightly larger which adds to the calibration complexity. However, we believe that by increasing the quality of the light source, these effects could be improved upon.

In general, we achieved higher quality projections with the rhombic prism where only two refractions occur and the refractive surfaces are perfectly parallel. The additional refractive surfaces in the two prism setup also introduce more stray light, though in both cases stray light is acceptable. A comparison of the same image projected with the rhombic prism and the two prism setup is shown in Figure 11. While the two-prism setup creates inferior results, we included it nonetheless as it shows how our method can be implemented with commercially available, mass-produced prisms.

With masks optimized towards light throughput (shown in Figure 13), the optimization reduces the color reproduction quality, though even for a maximum light throughput, the color representation is still acceptable. Since obfuscated masks are distributing the color information over all sub-pixels, the sub-pixel structure is not strongly apparent in the final projections.

The main source of errors in our physical results are misalignments of the two masks. In the simple case, the positioning need only be relative to the slit array of the first mask, pro-




Figure 11: The same images projected with the rhombic prism (left) and the two prism setup (right).




Figure 12: A series of experiments showing the original image input (left), the second prism mask (middle left), the simulated results (middle right) and a photograph of the actual projections (right).

viding more possibilities for a correct calibration. Optimized masks, on the other hand, need to be positioned absolutely. A horizontal positional misalignment causes a color shift in the projection similar to a hue shift since the position of the second mask defines which part of the spectrum is modulated. A rotational misalignment causes the color shift to vary over the projection (see Figure 14). The optimized masks cause

slightly worse artifacts since the projection content is defined by both masks.

Though we make sure that the masks are mounted flat on the prisms and are correctly aligned, we still see slight low-frequency color variations. We believe that the printed masks contain slight spatial drifts that are the result of the printing



Figure 13: A series of experiments showing the first masks (left), the second masks (middle left), simulated results (middle right) and a photograph of the actual projections with optimized masks (right). The top row applies masks that double the light throughput, the middle row triples the light output and the bottom row shows the result of obfuscated masks.



Figure 14: Misalignments of the masks compared to a proper alignment (left). A positional misalignment causes a color shift similar to a hue shift (middle). Rotational misalignments varies the hue shift over the projection (right).

process where the polymer foil is heated by the printer and later shrinks, introducing small deviations from the ideal mask.

6. Discussion and Future Work

Calibration & Error. Our relatively simple approach is limited by the resolution/accuracy of the printed masks and the practical accuracy in aligning the components. With custom prism design, higher-quality masks, and alignment equipment that allows for micrometer modifications, as well as an improved light source, we believe we could increase the resolution as well as the quality of the projections significantly.

Multi-Spectral Projection. Though our current prototypes focused on three color RGB projections, our setup actually allows us to trade color accuracy for light efficiency in a continuous fashion. We plan to explore multi-spectral projection capabilities by setting a smaller slit size than $e/2$ — trading light efficiency for spectral accuracy. With these features, combined with mask optimization, we believe our dispersion-based design could some day provide a compelling alternative

to traditional projection technology given professionally manufactured and calibrated components.

Dynamic Projection. Though we developed our prototypes as static/passive display devices using printed transparencies, we could easily use greyscale, high-resolution transmissive LCDs as the masks. This would achieve higher-quality, higher-resolution dynamic projections, and would allow for a dynamic selection of the light efficiency/color accuracy trade-off. Using greyscale LCDs would also eliminate the error introduced by binary dithering of the second mask.

Optimization & Cryptography. Though we have demonstrated increased light efficiency and obfuscation using optimized masks, more extensive analysis in this direction could prove fruitful. We believe it should be possible to optimally choose the trade-off between maximum color fidelity/saturation and light throughput by developing a more sophisticated, perceptual error metric. Furthermore, though we have demonstrated the potential for image obfuscation using optimized masks, we have not considered cryptographic strength and, unlike visual cryptography [NS95], it may be possible to infer the image from our individual masks. True cryptographic strength, however, could be trivially attained by splitting the second mask pattern into two separate masks that are superimposed, as in visual cryptography. A more interesting possibility would be to incorporate cryptographic guarantees directly into our optimization procedure.

7. Conclusion

We have described a simple and inexpensive color projection system which exploits dispersion as a source of color. Our prototype solely consists of one or two prisms and two monochromatic masks attached to the prism surfaces. The masks are produced with a standard Laser printer, achieving projections with a resolution of up to 100×100 pixel and making it easy to reproduce our results. We derived a simple analytic method to compute masks to project arbitrary color images and extended this idea with mask optimization to increase the light throughput (with approximations in the colors) and to obfuscate the image content.

We hope that our projection method poses new directions of projecting images, improving on light efficiency and spectral reproduction, acting as a building block in future developments.

References

- [AM10] ALEXA M., MATUSIK W.: Reliefs as images. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 4 (July 2010), 60:1–60:7. doi:10.1145/1778765.1778797. 1, 2
- [BBAM12] BERMANO A., BARAN I., ALEXA M., MATUSIK W.: Shadowpix: Multiple images from self shadowing. *Computer Graphics Forum (Proc. Eurographics)* 31, 2 (May 2012), 593–602. doi:10.1111/j.1467-8659.2012.03038.x. 1, 2
- [BKB*12] BARAN I., KELLER P., BRADLEY D., COROS S., JAROSZ W., NOWROUZEZHAI D., GROSS M.: Manufacturing layered attenuators for multiple prescribed shadow images. *Computer Graphics Forum (Proc. Eurographics)* 31, 2 (May 2012), 603–610. doi:10.1111/j.1467-8659.2012.03039.x. 1, 2
- [CDT*11] CAO X., DU H., TONG X., DAI Q., LIN S.: A prism-mask system for multispectral video acquisition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33 (2011), 2423–2435. doi:10.1109/TPAMI.2011.80. 2
- [Dua03] DUARTE F.: *Tunable Laser Optics*. Elsevier Science, 2003. 2
- [FDL10] FINCKH M., DAMMERTZ H., LENSCH H.: Geometry construction from caustic images. In *Proc. of European Conference on Computer Vision (ECCV)* (Sep 2010). doi:10.1007/978-3-642-15555-0_34. 1, 2
- [FRSL08] FUCHS M., RASKAR R., SEIDEL H.-P., LENSCH H. P. A.: Towards passive 6D reflectance field displays. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 27, 3 (2008), 58:1–58:8. doi:10.1145/1399504.1360657. 2
- [KSIN07] KASAROVA S. N., SULTANOVA N. G., IVANOV C. D., NIKOLOV I. D.: Analysis of the dispersion of optical plastic materials. *Optical Materials* 29, 11 (2007), 1481–1490. doi:10.1016/j.optmat.2006.07.010. 3
- [LWLD14] LIN X., WETZSTEIN G., LIU Y., DAI Q.: Dual-coded compressive hyperspectral imaging. *Opt. Lett.* 39, 7 (Apr 2014), 2044–2047. URL: <http://ol.osa.org/abstract.cfm?URI=ol-39-7-2044>, doi:10.1364/OL.39.002044. 2
- [MP09] MITRA N. J., PAULY M.: Shadow art. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 28, 5 (2009). URL: 10.1145/1661412.1618502. 1, 2
- [MRT08] MOHAN A., RASKAR R., TUMBLIN J.: Agile spectrum imaging: Programmable wavelength modulation for cameras and projectors. *Computer Graphics Forum (Proc. Eurographics)* 27, 2 (2008), 709–717. doi:10.1111/j.1467-8659.2012.03036.x. 2
- [NS95] NAOR M., SHAMIR A.: Visual cryptography. In *Advances in Cryptology — EUROCRYPT '94*, vol. 950 of *Lecture Notes in Computer Science*. Springer Berlin/Heidelberg, 1995, pp. 1–12. doi:10.1007/BFb0053419. 2, 10
- [PJJ*11] PAPAS M., JAROSZ W., JAKOB W., RUSINKIEWICZ S., MATUSIK W., WEYRICH T.: Goal-based caustics. *Computer Graphics Forum (Proc. Eurographics)* 30, 2 (June 2011), 503–511. doi:10.1111/j.1467-8659.2011.01876.x. 1, 2
- [WDB*07] WEYRICH T., DENG J., BARNES C., RUSINKIEWICZ S., FINKELSTEIN A.: Digital bas-relief from 3D scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 26, 3 (July 2007), 32:1–32:7. doi:10.1145/1276377.1276417. 2
- [WLHR11] WETZSTEIN G., LANMAN D., HEIDRICH W., RASKAR R.: Layered 3D: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Transactions Graphics (Proc. SIGGRAPH)* 30, 4 (2011). doi:10.1145/2010324.1964990. 1, 2
- [WPMR09] WEYRICH T., PEERS P., MATUSIK W., RUSINKIEWICZ S.: Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 28, 3 (July 2009), 32:1–32:6. doi:10.1145/1531326.1531338. 2
- [YIC*12] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Pixel art with refracted light by rearrangeable sticks. *Computer Graphics Forum* 31, 2 (May 2012), 575–582. doi:10.1111/j.1467-8659.2012.03036.x. 2