

An Analysis of Spatial- and Fourier-Multiplexed Imaging

Gordon Wetzstein

Ivo Ihrke
The University of British Columbia

Wolfgang Heidrich

Technical Report
UBC CS TR-2009-06

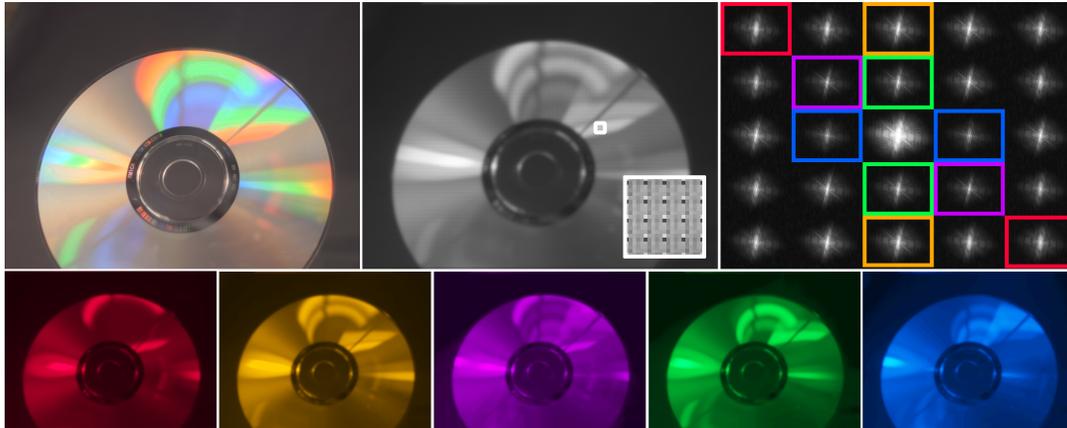


Figure 1: A Fourier-multiplexing approach to multi-spectral imaging. Upper row, from left: focused color picture of a CD, a monochromatic sensor image of a band-limited scene shot through our multi-spectral mask, Fourier transform of the sensor image. Lower row: 5 of the 24 reconstructed color channels.

Abstract

Multiplexing is a common technique for encoding high-dimensional image data into a single two-dimensional image. Examples of spatial multiplexing include Bayer patterns to encode color channels, and integral images to encode light fields. In the Fourier domain, optical heterodyning has been used to encode light fields.

In this paper, we analyze the relationship between spatial and Fourier multiplexing techniques. We develop this analysis on the example of multi-spectral imaging, and then generalize it to light fields and other properties. We also analyze the effects of sensor saturation on Fourier multiplexing techniques such as optical heterodyning, and devise a new optimization approach for recovering saturated detail.

CR Categories: I.4.1 [IMAGE PROCESSING AND COMPUTER VISION]: Digitization and Image Capture;

1 Introduction

Two-dimensional images are the primary means by which humans represent the three-dimensional world surrounding them. The introduction of photography resulted in unprecedented levels of realism in these representations. From the earliest works of photographing through an array of pinholes at the beginning of the last century, to current light field cameras, an often-stated goal in photography has been to increase the flexibility of the image capturing process. Simultaneous acquisition of several views on a single image plane is just one example of these developments.

In fact it has recently been argued that the “ultimate” camera would capture, within the limits of the uncertainty principle, the position, direction, wavelength, and time of arrival of each individual photon incident on the image sensor [Tumblin 2007]. That is, the ideal camera would be able to capture a 2D slice of the plenoptic function [Adelson and Wang 1992]. From this information it would then be possible to extract different styles of images using computation. A practical challenge to achieving this goal is that real world image sensors always integrate over a finite area, incident angle, wavelength, and time. This makes it difficult to use a two-dimensional sensor to capture simultaneous detail in more than two of the mentioned dimensions.

A common solution to this problem is to trade spatial resolution for an encoding of additional dimensions into the image plane. This *multiplexing* can be achieved with something as simple as a color filter array like the Bayer pattern [1976], or a more complicated encoding of multiple dimensions using Assorted Pixels [Nayar and Narasimhan 2002; Narasimhan and Nayar 2005]. A spatial en-

coding for light fields can be achieved by multiplexing different views into different spatial frequency bands using optical heterodyning [Veeraraghavan et al. 2007].

In this paper, we analyze the relationship between spatial and Fourier multiplexing techniques. We develop this analysis on the example of multi-spectral imaging, for which we introduce a new Fourier-multiplexed formulation. These observations translate to other image multiplexing approaches, including integral and light field photography.

We also show that all existing Fourier multiplexing approaches are susceptible to severe artifacts in the case of sensor saturation. We analyze these artifacts, and develop a novel optimization approach for recovering the signal in saturated regions. The result is a Fourier-space reconstruction algorithm that recovers a larger dynamic range than spatial algorithms using the same filter patterns.

To summarize, our technical contributions are

- a theoretical framework comparing spatial and Fourier multiplexed imaging,
- an analysis of saturation artifacts in Fourier multiplexing,
- a novel Fourier multiplexing approach for multi-spectral imaging, and
- a new HDR reconstruction algorithm operating in Fourier space.

2 Background and Related Work

As indicated in the previous section, physically realizable image sensors integrate over a finite range along each dimension of the plenoptic function. In order to capture multiple samples along any of the non-spatial dimensions, three fundamental approaches are available.

Multi-sensor capture refers to an approach where multiple image sensors simultaneously capture different samples of the plenoptic function. Examples of this approach include multi-camera systems for capturing light fields and related information, as well as 3-chip cameras, which have a separate image sensor for the red, green, and blue channel of a color image. A similar approach is available for HDR imaging [Aggarwal and Ahuja 2004; McGuire et al. 2007].

Time sequential capture can be used when the hardware requirements of multi-sensor capture would be prohibitive. For example, light fields can also be captured by moving a single camera to different positions [Levoy and Hanrahan 1996; Gortler et al. 1996]. Color images can be acquired sequentially by applying different color filters (e.g. [Wang and Heidrich 2004]), which is particularly attractive when a large number of color channels is desired, for example in microscopy. HDR images can be acquired by sequentially capturing different exposures [Debevec and Malik 1997; Mitsunaga and Nayar 1999], or using generalized image mosaicing [Schechner and Nayar 2003].

A downside of time-sequential capture is of course the difficulty of capturing dynamic environments and videos. If neither multi-sensor nor time sequential capture are feasible, one can employ the third approach:

Multiplexed imaging refers to the multiplexing of multiple images onto a single image/sensor plane. In effect, this approach trades spatial resolution for the ability to simultaneously capture

multiple slices of the plenoptic function. Multiplexing can either be performed in the spatial domain or in the Fourier domain.

In spatial multiplexing, the pixels in a local neighborhood of the image sensor represent samples from different slices of the plenoptic function. Most often, the image is comprised of repeating tiles, which we call *super-pixels*. Corresponding pixels in each super-pixel come from the same slice of the plenoptic function. The most commonly used example of spatial image multiplexing is the use of a color filter array such as the Bayer pattern [Bayer 1976] in order to acquire a color image in a single photograph. General color filter arrays can be used to capture high dynamic range or multi-spectral images [Nayar and Narasimhan 2002; Narasimhan and Nayar 2005]. Integral and light field photography can be seen as multiplexing of the light field onto a 2D image plane [Lippmann 1908; Adelson and Wang 1992; Ng 2005]. Here, the super-pixels correspond to different viewpoints, while corresponding points within each super-pixel correspond to the same direction.

Fourier-space multiplexing is very closely related to spatial multiplexing, but here the spatial structure is optimized to achieve certain properties in the Fourier space. In particular, spatial patterns are chosen such that the different slices of the plenoptic function are encoded into different frequency bands. In computer graphics, this *optical heterodyning* approach has so far been used for capturing light fields [Veeraraghavan et al. 2007; Veeraraghavan et al. 2008] as well as occluder information [Lanman et al. 2008]. Georgiev et al. [2008] recently analyzed spatially encoded light fields in Fourier space and demonstrated that similar reconstruction algorithms apply here. In imaging, a similar approach has been used to analyze filter patterns for Bayer-style color mosaics [Alleyson et al. 2005].

It is important to point out that both image-space and Fourier-space multiplexing require the image to be spatially low-pass filtered to the super-pixel resolution to avoid aliasing. For example, digital cameras commonly employ an *anti-aliasing filter*, which blurs the image to the resolution of a super-pixel in the Bayer pattern (see e.g. [Greivenkamp 1990]). The camera software later partially compensates for this blur by sharpening the raw image.

Various other methods have been proposed that employ combined mask or lens-based optical light modulation and computational reconstruction. These approaches demonstrate how to capture images and reconstruct scene depth [Levin et al. 2007], remove motion blur from a single photograph [Raskar et al. 2006], or remove veiling glare from photographs [Raskar et al. 2008; Talvala et al. 2007].

3 Overview

Optical heterodyning as introduced by Veeraraghavan et al. [2007] multiplexes a 4D light field into a 2D image. The 2D Fourier transform of the image contains tiles that can be re-arranged into a 4D Fourier “volume”. The actual light field can then be reconstructed from this 4D function via the inverse Fourier transform.

In the following, we demonstrate how this approach can be generalized to other dimensions of the plenoptic function (Section 4.1). We will use the specific case of multi-spectral imaging, in which the tiles of the Fourier transform represent a 3D spatio-chromatic volume that can be recovered via a 3D inverse Fourier transform.

A direct application of this heterodyning approach results in color filter arrays that are difficult to manufacture and have low transmissivity. We show how the approach can be extended to the multiplexing of arbitrary spectral basis functions, which can then be produced more easily, and are significantly more transmissive (Section 4.2).

We demonstrate the approach using a scan camera prototype (Section 4.3).

In Section 5 we then analyze the relationship between this new Fourier multiplexing approach and existing spatial color mosaics such as the Bayer pattern [1976] and Assorted Pixels [Nayar and Narasimhan 2002]. We demonstrate that it is possible to reconstruct the newly derived patterns with both (local) spatial methods as well as (global) Fourier reconstruction methods. We also show that the same is true for traditional Bayer patterns and related mosaics. These findings also translate to encodings of other properties of the plenoptic function. We show that light field photography images (e.g. [Ng 2005; Levoy et al. 2006]) can be reconstructed with (global) Fourier algorithms, and Dappled Photography images [Veeraraghavan et al. 2007] can be reconstructed using (local) spatial approaches.

In Section 6, we discuss the problem of sensor saturation, and show it to be a fundamental problem in all Fourier multiplexed approaches. We discuss how this problem can be solved using an optimization-based approach (Section 7), which allows us to capture HDR images.

4 Multi-Spectral Fourier Multiplexing

In this section we discuss concepts that extend optical heterodyning to multiplexed spectral imaging. We present a framework that generalizes grayscale heterodyne masks and corresponding computational reconstruction algorithms to encodings of the color spectrum. By exploiting the separability of certain dimensions of the plenoptic function for multiplexed acquisition in the Fourier domain, we show how to design color masks that can easily be manufactured, and have a superior light transmissivity compared to both standard color filters, as well as grayscale heterodyne masks.

4.1 Multi-Spectral Heterodyning

For Fourier multiplexing of multi-spectral images, we would like to work with a 3D *spatio-chromatic* space, consisting of two spatial dimensions x and y , and the wavelength of light λ . A heterodyne mask $M(x, y, \lambda)$ embeds slices of the Fourier transform of this 3D space as tiles into the 2D Fourier transform of a multiplexed image. In the case of light field multiplexing with Dappled Photography [Veeraraghavan et al. 2007], the heterodyning is achieved by a spatially varying neutral density filter that is offset from the image sensor, and therefore provides an encoding of parallax (view directions). In the case of multi-spectral imaging, we use a filter that is located *in* the sensor plane and has spatially varying spectral transmission characteristics. The N spectral slices are encoded into $N_x \times N_y$ Fourier tiles

$$\begin{aligned}
 M(x, y, \lambda) &= \\
 &= \sum_{k=-p_x}^{p_x} \sum_{l=-p_y}^{p_y} \cos(2\pi(kf_x^0 x + lf_y^0 y + (l \cdot N_x + k)f_\lambda^0 \lambda)) \quad (1) \\
 &= \sum_{k,l} \cos(2\pi(kf_x^0 x + lf_y^0 y)) \cdot \cos((l \cdot N_x + k)f_\lambda^0 \lambda) \\
 &\quad + \sum_{k,l} \sin(2\pi(kf_x^0 x + lf_y^0 y)) \cdot \sin((l \cdot N_x + k)f_\lambda^0 \lambda), \quad (2)
 \end{aligned}$$

where $N_x = 2 \cdot p_x + 1$, and $N_y = 2 \cdot p_y + 1$. A similar relationship holds if an even number of tiles is being encoded in x or y .

This approach also generalizes to other dimensions of the plenoptic function. For example, temporal information could be encoded using an in-plane filter that varies both in space and in time, i.e. a partially transparent display.

When capturing a spatially band-limited scene through the filter from Equation 1, spectral information is optically transformed into spatial frequencies. Specifically, the mask creates multiple copies of the scene around the fundamental spatial frequencies f_x^0 and f_y^0 in the Fourier transform of the captured image, which contains $N = N_x \times N_y$ copies in total. The wavelength dependent term in the mask enables the sensor to sample three-dimensional Fourier coefficients with a distance of f_λ^0 in the wavelength spectrum. The color spectrum can be directly extracted from the Fourier transform of the acquired image by cropping the individual tiles, stacking them, and performing a multi-dimensional inverse Fourier transform.

Equation 2 demonstrates that the channels are multiplexed with both the sine and the cosine basis functions. Specifically, the cosines in the spectral dimension are encoded in the cosines of the spatial Fourier basis. Likewise, the spectral sines are encoded in the sines of the spatial Fourier basis. Together, the spatial Fourier basis encodes N different spectral basis functions, which happen to also represent a Fourier basis in the spectral dimension.

The center row of Figure 2 shows a 3×3 color pattern and the transmission spectra of 3 of its pixels. A number of things can be noted here. First, very specific transmission curves are required for each of the pixels. These may be difficult to manufacture with sufficient accuracy. Second, the spectra have relatively well-defined maxima, resulting in dark filters (although brighter than a corresponding Assorted Pixels pattern, shown in the top row).

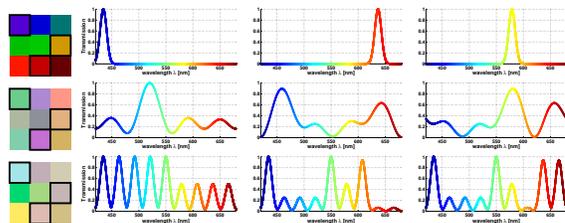


Figure 2: The repetitive patterns that make up different multi-spectral masks converted to sRGB (left column from top: assorted pixels, spectral heterodyne, and arbitrary spectral bases mask). Columns 2 to 4 show spectral plots of different pixels in the individual masks (highlighted on the left).

4.2 General Multiplexing of Spectral Basis Functions

Standard color printing and digital film transparency exposure techniques limit the spectral complexity of color masks that can be manufactured. Masks that are built using these techniques can only be composited of spatially varying combinations of spectral distributions, which are predefined by the color channels of the printing device. With these restrictions, color masks with a very complex spectral shape cannot be produced. To address this issue, we can use arbitrary repetitive patterns that have an impulse train as their multi-dimensional Fourier transform [Lanman et al. 2008].

In particular, we can replace the spectral Fourier basis in Equation 2 with arbitrary basis spectra $b_{k,l}^{c,s}(\lambda)$ corresponding to the primary colors of a printing device:

$$M(x, y, \lambda) = \sum_{k,l} \cos(2\pi(kf_x^0 x + lf_y^0 y)) b_{k,l}^c(\lambda) + \sum_{k,l} \sin(2\pi(kf_x^0 x + lf_y^0 y)) b_{k,l}^s(\lambda). \quad (3)$$

Note that we still use a Fourier basis for the spatial dimensions, so that the symmetries are the same as before, with $b_{k,l}^c = b_{-k,-l}^c$ and a similar relationship for the sine basis. In total, there are $N/2$ distinct bases $b_{k,l}^c$ encoded in the sum of cosines, with the remaining number of distinct bases $b_{k,l}^s$ being encoded in the sum of sines.

The result of this encoding is that the 2D Fourier transform of the multiplexed image L_s contains tiles $T_{k,l}$ that correspond to the 2D Fourier transform of the original signal, filtered by a specific spectral distribution as determined by the spectral basis functions. The specific symmetries of the spatial Fourier transform mean that a pair of two of these tiles $T_{k,l}$ and $T_{-k,-l}$ together encodes two basis functions $b_{k,l}^c$ and $b_{k,l}^s$. This relationship is detailed in Appendix A.

Hence, after decomposition, each channel can be reconstructed by cropping the corresponding Fourier tiles and performing a two-dimensional inverse Fourier transform. This process is illustrated in Figures 1 and 5.

Filter Normalization. When implementing the filter patterns discussed above, an obvious constraint is that the transmission has to be in the range $[0 \dots 1]$ for each wavelength. That is, the mask can neither amplify a wavelength, nor produce negative light. In practice, the constraints may be more stringent, for example due to limited contrast in the film material. These constraints can easily be incorporated by renormalizing the individual mask coefficients for each color channel.

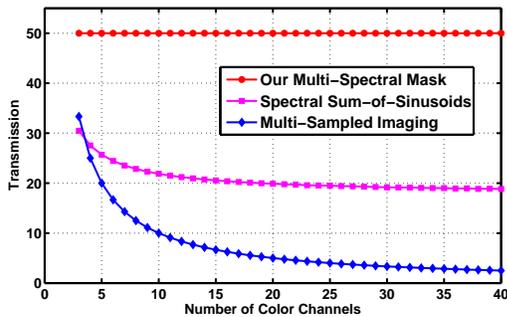


Figure 3: Comparison of spectral mask light transmittances for multi-sampled imaging, spectral sum-of-sinusoids and our mask.

When comparing the new mask with the heterodyne mask from Section 4.1, we see that the transmissivity is significantly increased, especially for large numbers of bases (Figure 3). The integral of a single normalized sinusoid is 50%. The total transmissivity of our mask is, therefore, half of the sum of the individual transmissivities for the individual primaries b_k .

4.3 Proof of Concept Implementation

Simulation. To validate the concepts with a large number of color channels, we used a partial simulation approach, in which a spatially varying filter is simulated by sequential exposures with different color filters. This partial simulation provides real multi-spectral

data with real camera noise without the need to manufacture a spatially varying filter with many primaries.

For Figure 1, we captured 12 multi-exposure sequences with uniform color filters. Twelve additional color channels were synthetically generated as linear combinations of the captured ones. The band-limitation was achieved by defocusing the camera and slightly blurring the images. We then synthetically combined pixels from the different exposure sequences to generate a synthetic multiplexed image using the normalized version of our multi-spectral mask described in Equation 3. The reconstruction, also depicted in Figure 1, illustrates that the channels can be faithfully reconstructed from a single sensor image.

Prototype. A full prototype with a more limited number of color channels was constructed using a large format camera and a flatbed scanner [Wang and Heidrich 2004]. This camera is easy to build, provides a very high-resolution, and is large-scale, which makes it easy to use simple transparencies as masks. We attached our multi-spectral masks directly onto the glass plate, where the incoming light as well as the underlying sensor elements are focused. A holographic diffuser, which allows the incident light to form an image, is also mounted over the filter on the scanner glass plate as seen in Figure 4.



Figure 4: Our prototype is a large format high-resolution scanner camera. The filters are mounted under a diffuser on the scanner's glass plate.

Our masks are high resolution RGB digital images, exposed onto photographic film using light-valve technology (LVT). Color transparencies with a resolution of up to 2032 dpi and a high contrast can be ordered at professional print service providers such as Bowhaus (www.bowhaus.com). We used 4"x5" transparencies and performed scans with 2400 dpi, which results in fairly high resolution images.

Figure 5 shows results that we have acquired using our prototype, and the RGB mask we used (inset on the upper left). The Fourier transform (center left) of the captured grayscale image (upper left) contains three copies of the original signal. Each one is filtered through one of the color filters and the DC tile contains the sum of all color channels. The tiles can be cropped and individually inverse Fourier transformed to recover the spatial channels (lower left).

The scan camera prototype suffers from high noise and stripe artifacts. However, we observe that the DC tile has a higher signal level, and thus a better signal-to-noise ratio, than the other Fourier tiles. This is due to the normalization of the pattern to positive transmission values. We can use this fact to improve the reconstruction as follows. We transform the image into a YUV-like space, where we have an intensity channel in which all color channels are weighed equally. We then replace the Y channel with the reconstructed DC Fourier tile, and transform the image back to RGB

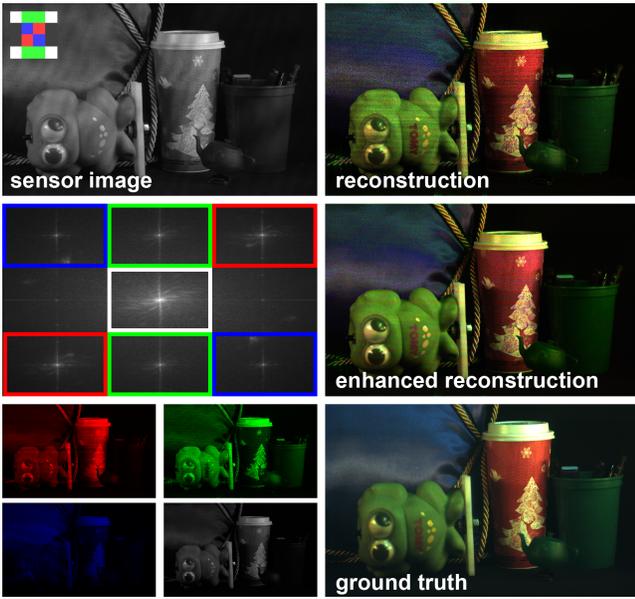


Figure 5: Multi-spectral results acquired with an RGB mask and our scanner camera. The left column shows the captured grayscale image with an inset of the employed color mask, the Fourier transform of the sensor image, and the colored reconstructed channels. Noise in the reconstructed color image (upper right) can be reduced using our image enhancement technique (center right).

space (Figure 5 center right).

For verification we have captured three images of the same scene through transparencies of the individual color channels (Figure 5 lower right). Slight intensity changes are due to vignetting, because ground truth and mask modulated image are captured on different areas on the scanner surface.

5 Analysis of Multiplexed Imaging

After introducing a new Fourier multiplexing approach for multi-spectral imaging, we now discuss the relationship between spatial and Fourier multiplexing. This relationship holds for multi-spectral imaging as well as multiplexing of the other dimensions of the plenoptic function (direction and time).

For this purpose we introduce a general Fourier-multiplexing mask for attributes of the plenoptic function. The key criterion for a mask to be a Fourier multiplexing mask is that it modulates the signal with a discrete number of spatial frequencies. This insight is a motivation to split the mask into two separable parts, one part describing the *plenoptic basis*, and one part for the *spatial basis*, which is the Fourier basis:

$$\begin{aligned}
 M(x, y, \vec{\omega}, \lambda, t) &= \sum_{k,l} b_{k,l}(\vec{\omega}, \lambda, t) \cos(2\pi(kf_x^0 x + lf_y^0 y)) \\
 &= \sum_{k,l} b_{k,l}(\vec{\omega}, \lambda, t) s_{k,l}(x, y). \quad (4)
 \end{aligned}$$

For simplicity of notation we only consider the cosine part of the spatial encoding. The formula including the sine part can be derived in a similar manner. The plenoptic basis is described by $N_x \times N_y$ different basis functions $b_{k,l}$. The $N_x \times N_y$ spatial frequencies are

used to encode the plenoptic basis into the spatial basis $s_{k,l}$. Now consider the image L_s formed by some sensor through this type of mask

$$\begin{aligned}
 L_s(x, y) &= \iiint_{\Omega, \lambda, t} L(x, y, \vec{\omega}, \lambda, t) M(x, y, \vec{\omega}, \lambda, t) d\vec{\omega} d\lambda dt \\
 &= \sum_{k,l} s_{k,l}(x, y) \iiint_{\Omega, \lambda, t} L(x, y, \vec{\omega}, \lambda, t) b_{k,l}(\vec{\omega}, \lambda, t) d\vec{\omega} d\lambda dt \\
 &= \sum_{k,l} c_{k,l}(x, y) s_{k,l}(x, y). \quad (5)
 \end{aligned}$$

We see that the projection into the plenoptic basis is performed by the integrating properties of the sensor, yielding the (dual) plenoptic basis coefficients $c_{k,l}(x, y)$ multiplying the spatial basis $s_{k,l}(x, y)$. Note that the plenoptic coefficients vary with the spatial sensor location.

Relationship of Spatial and Fourier Reconstruction. Both spatial and Fourier multiplexing of plenoptic information require a trade-off between the spatial resolution and the number of additional channels being recorded. This trade-off requires band-limiting of the incident information on the image sensor. If the image is not properly band-limited, spatial sampling approaches violate the sampling theorem, whereas Fourier-multiplexing techniques will observe overlap in the frequency bands occupied by different parts of the signal. Both conditions are related by the duality between the spatial and the Fourier representation of the signal. To avoid aliasing, an optical anti-aliasing filter is usually mounted in front of camera sensors [Greivenkamp 1990].

Spatial Reconstruction of Fourier-Multiplexed Images. By capturing information with Fourier-multiplexing masks using a number of plenoptic basis functions, weighted combinations of those basis functions are being recorded by the sensor [Veeraraghavan et al. 2007; Lanman et al. 2008]. The de-multiplexing operation is the Fourier transform performed on a global scale, i.e. on the whole image. We now show that *spatial reconstruction*, that is, reconstruction using linear combinations of the recorded signal can be performed using only *local* information from the multiplexed image.

For clarity we are restricting ourselves to the analysis of repeating image space patterns, which we call *super-pixels*. Lanman and co-workers [2008] have shown that such repeating patterns give rise to a series of Dirac peaks in the Fourier domain, and are thus generating the spatial basis $s_{k,l}$ as required. The optical low-pass filtering ensures that the radiance incident at the filter pattern varies smoothly at the scale of a super-pixel.

Proposition. *The local linear transformation that recovers the plenoptic basis coefficients from a suitably band-limited Fourier-multiplexed image is the inverse spatial transformation performed at the scale of a super-pixel.*

Proof. To prove this property, we re-consider Equation 5. Introducing the (integer) fundamental frequencies $f_x^0 = s_x/N_x$ and $f_y^0 = s_y/N_y$, where (s_x, s_y) is the sensor resolution, we analyze this equation at the resolution of one super-pixel. At this resolution, both f_x^0 and f_y^0 are equal to one. We now apply the inverse spatial transformation (i.e. the Fourier transform) within a super-pixel, giving

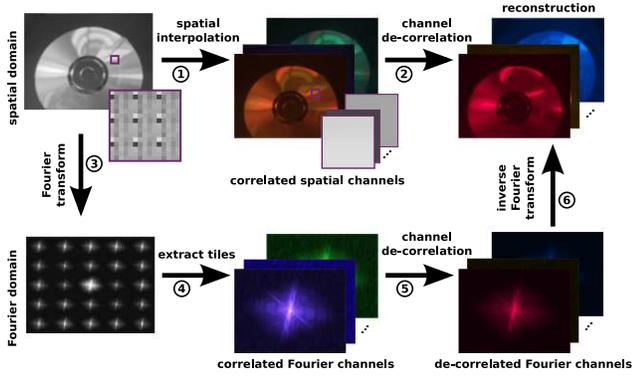


Figure 6: An illustration of the duality of spatial and Fourier-space de-multiplexing.

$$\begin{aligned}
 F\{L_s(x, y)\} &= F\left\{\sum_{k,l} c_{k,l}(x, y) \cdot \cos(2\pi(kx + ly))\right\} \\
 &= \sum_{k,l} \frac{1}{2} \delta(f_x + k, f_y + l) \otimes F\{c_{k,l}(x, y)\} + \\
 &\quad \frac{1}{2} \delta(f_x - k, f_y - l) \otimes F\{c_{k,l}(x, y)\}. \quad (6)
 \end{aligned}$$

We see that every Fourier coefficient of the super-pixel is a sum of the Fourier transformed plenoptic coefficients $c_{k,l}$. However, if the DC component of the Fourier transformed plenoptic coefficients is the dominating term, that is, if the other plenoptic Fourier coefficients are negligible, then the convolution reduces to a multiplication by a constant:

$$\begin{aligned}
 F\{L_s(x, y)\} &\approx \sum_{k,l} \frac{1}{2} \delta(f_x + k, f_y + l) \cdot c_{k,l}(x, y) + \\
 &\quad \frac{1}{2} \delta(f_x - k, f_y - l) \cdot c_{k,l}(x, y). \quad (7)
 \end{aligned}$$

Now we see that each Fourier coefficient of a super-pixel corresponds to a close approximation of exactly one plenoptic basis coefficient. The condition that the DC coefficient of $F\{c_{k,l}(x, y)\}$ be dominant is what we refer to as a suitably band-limited signal. \square

Although the image is band-limited, the higher frequencies are in practice not completely negligible over the extended region of a super-pixel. However, the band-limit implies that the spatial variations are of a low-enough frequency that it is possible to interpolate corresponding channels of the plenoptic function from neighboring super-pixels (Figure 6 (1)). This interpolation directly corresponds to spatial de-mosaicing algorithms used in cameras with color filter arrays such as the Bayer pattern.

In general, the recovered plenoptic coefficients $c_{k,l}$ are correlated by a matrix that encodes linear combinations of the channels for a particular pixel. This matrix must be inverted (Figure 6 (2)) to obtain channels that represent coefficients of the original basis $b_{k,l}$.

Fourier Reconstruction of Spatially Multiplexed Images. Just as Fourier-multiplexed images can be reconstructed spatially, we can also Fourier-reconstruct spatially multiplexed images. The

Fourier view helps, for example, to analyze aliasing artifacts and other issues. It also leads to a better understanding of the properties of plenoptic sampling and can be used advantageously e.g. in de-mosaicing of color filter arrays [Dubois 2005]. Because of limited space we omit a formal proof, but give an intuitive explanation.

By again analyzing the spatial sampling as a masking operation, we see that each individual sensor element records a different plenoptic basis coefficient. Because the plenoptic filters are repeating patterns, they form Dirac peaks in Fourier space (Figure 6 (3)). These tiles can be cut out and stacked into multiple channels (Figure 6 (4)). However, the Dirac peaks superimpose, and weighted averages of the plenoptic basis coefficients are sampled in the Fourier domain. Thus, the channels in the stacked image are, in general, correlated by a matrix. This matrix needs to be inverted (Figure 6 (5)) to obtain the Fourier representation of the plenoptic basis coefficients. Finally, the spatial-domain reconstruction can be obtained through an inverse Fourier transform of the individual channels (Figure 6 (6)).

Interpretation. The previous two theorems demonstrate that the difference between spatial and Fourier multiplexing lies primarily in the reconstruction algorithm. The patterns themselves can be used with either method. However, as we shall see in the examples below, patterns developed for Fourier multiplexing tend to produce uncorrelated channels in Fourier space, so that the matrix in Figure 6(5) is the identity matrix. Likewise, patterns developed for spatial multiplexing, tend to produce (almost) uncorrelated samples in the spatial domain, so that the matrix (2) is close to the identity matrix. In general, however, one can design patterns for special purposes that are correlated in both the Fourier and the spatial domain. The Fourier-space view can help optimizing patterns for special properties, such as better transmissivity in the case of multi-spectral imaging.

Although we have demonstrated a duality between spatial and Fourier-space reconstruction, this result should not be interpreted as meaning that both algorithms produce the exact same reconstruction results. Differences do arise in the way the individual channels are interpolated to the full image resolution. In the Fourier-space reconstruction, each pixel in the final result is a linear combination of the multiplexed image. In spatial reconstruction, other types of interpolation are frequently used, including edge-preserving interpolation [Chang et al. 1999], or more sophisticated methods such as learning-based techniques [Nayar and Narasimhan 2002]. In the following we show a number of examples of the duality between spatial and Fourier reconstruction.

Example: Spatial Reconstruction of Multi-Spectral Images.

We now re-visit the Fourier-multiplexed multi-spectral imaging example introduced in Section 4. The plenoptic basis is chosen as $b_{k,l}(\vec{\omega}, \lambda, t) = b_{k,l}(\lambda)$, i.e. a purely spectral basis. For spatial encoding we use the complete set of Fourier basis functions, i.e. cosines and sines. After application of the Fourier transform on a per super-pixel level, we have to perform either a 1D inverse Fourier transform for heterodyne multi-spectral imaging (Section 4.1), or a decomposition as detailed in Appendix A in the case of arbitrary basis functions (Section 4.2). The results of this experiment are shown in Figure 7.

Example: Spatial Reconstruction of Heterodyne Light Fields.

As a second example we consider light field imaging by optical heterodyning, as introduced by Veeraraghavan et al. [2007]. In this case, the plenoptic basis is purely directional. In their paper, the authors show that the directional light field component is expanded in the Fourier basis. That is, the Fourier copies of the multiplexed sig-

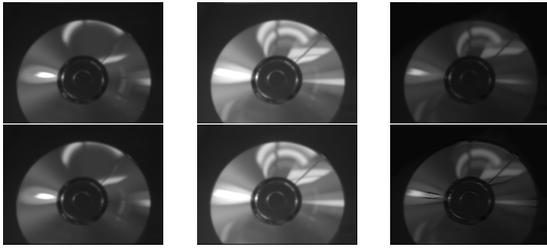


Figure 7: Fourier-based (top row) and spatial (bottom row) reconstruction of our Fourier-multiplexed multi-spectral dataset, channels 1, 4 and 8, respectively. Channel 8 is neighboring the DC copy in Fourier space and thus exhibits most aliasing effects from insufficient band-limitation. In this channel, differences between the two reconstruction approaches, primarily at sharp intensity transitions, can be observed.

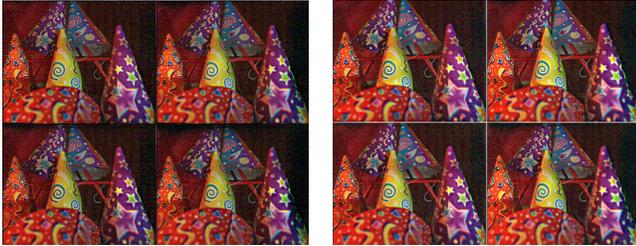


Figure 8: Fourier-based (left) and spatial reconstruction (right) of 2×2 views of the “Cones” data set of Veeraraghavan et al. [2007].

nal store the Fourier coefficients of the directional light field component. The derivation of the plenoptic and spatial basis separation is similar to the case of multi-spectral heterodyning (Equation 2). Using the results from our analysis, we see that we have to invert the spatial transformation by applying the Fourier transform on a super-pixel level to reconstruct its directional Fourier coefficients. If the sum-of-sinusoids mask employed by Veeraraghavan et al. [2007] was un-normalized, the inverse transform would reconstruct the an identical image, i.e. a pinhole view. However, since the mask is normalized to meet physical constraints, a re-weighting of the Fourier coefficients must be performed prior to performing the inverse Fourier transform. This corresponds to a high-pass filter performed on the super-pixel. An example of spatial reconstruction of the “Cones” data set is shown in Figure 8 (right). The reconstruction approach directly carries over to the MURA patterns employed by Lanman et al. [2008] for lens-less occluder imaging.

Example: Fourier Reconstruction of Bayer Patterns. By exploiting the duality between spatial and Fourier filtering approaches, we can now have a look at traditional color multiplex-

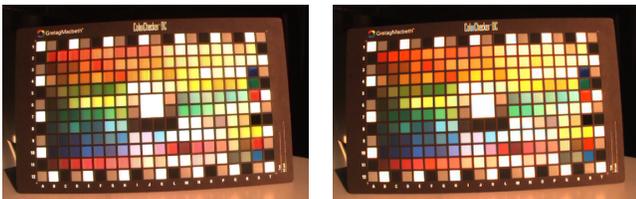


Figure 9: Bayer reconstruction using spatial interpolation (left) and by re-combination of Fourier-tiles (right).

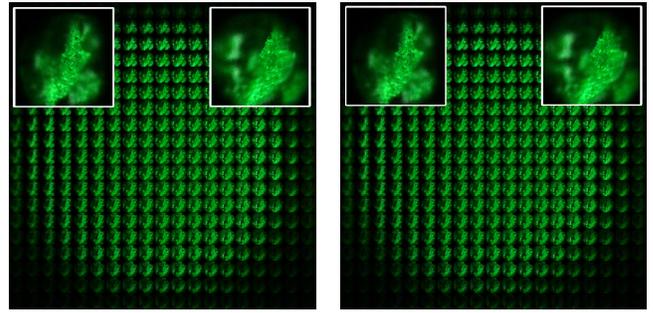


Figure 10: Light field reconstruction using spatial interpolation (left) and by re-combination of Fourier-tiles (right). The light field is property of Levoy et al. [2006].

ing techniques. Typically, color filters are arranged in a Bayer pattern [1976]. The spectral filter response for every red, green and blue pixel is measured directly. As we have seen before, the Fourier representations of the red, green and blue channels will be correlated in the Fourier copies created by a 2×2 grid of Dirac peaks. For the particular case of the Bayer pattern, it has been noted that the Fourier copies store luminance and chrominance, respectively [Ailleyson et al. 2005]. The DC peak is creating a monochromatic copy, whereas the other peaks record $2G - R - B$, $R - B$, and $B - R$ components, respectively. As we can see, only two of the three chromaticity components are independent. This is to be expected since there are only 3 color channels but 4 Fourier copies. For demonstration purposes we show spatial and Fourier-based demosaicing of a Macbeth DC ColorChecker in Figure 9. Note again that the results are the same up to pixel level interpolation.

Example: Fourier Reconstruction of Lenslet Array Images.

Our last example is Fourier-based reconstruction of light field data recorded with lenslet arrays. Georgiev et al. [2008] have already shown that such data can be reconstructed in Fourier space. In this case, the plenoptic basis is again directional, recording a different direction in every pixel of the spatially repeating pattern. A super-pixel is now the image of one lenslet. Each pixel within the super-pixel records exactly one direction. Again, in Fourier space, a number of copies is created that equals the number of pixels in the super-pixel. Each Fourier copy is a weighted average of the directional coefficients’ Fourier representations. In Figure 10 we show a comparison of spatial reconstruction by re-sampling the lenslet images, and the corresponding Fourier reconstruction. The latter was generated by de-correlating the Fourier copies prior to transforming back to spatial domain. The de-correlation is achieved by an inverse two-dimensional Fourier transform, and can be implemented by re-stacking the Fourier tiles of the original image into a 4D Fourier representation of the light field, followed by an inverse multi-dimensional Fourier transform. This is equivalent to the reconstruction proposed in [Veeraraghavan et al. 2007] without peak re-normalization.

6 Saturation Analysis in Fourier Space

We now turn our attention to the analysis of saturation effects in recorded imagery. Saturation in the presence of a high dynamic range scene is a severe limitation of most image sensors. In spatial reconstruction techniques, the effects of saturation are well known and easily understood. However, the problem has been ignored so far for Fourier-multiplexing approaches to image acquisition. In this section, we analyze this problem and identify a fundamental issue with Fourier-multiplexing techniques, for which we then prop-

pose a solution in Section 7.

Figure 11 shows a band-limited, one-dimensional signal and its clipped counter-part, as well as logarithmic plots of their Fourier transforms. The effects of clipping are clearly visible in the Fourier representation of the saturated signal. While the original signal is perfectly band-limited, the clipped signal exhibits significant high frequency components. This can be explained by sharp corners that are being introduced by clipping the smooth signal.

Note however, that the high frequency components introduced by saturation are not random, but are closely related to the missing data. This can be seen by decomposing the original (unsaturated) signal into the measured image and an error: $L = L_{\text{meas}} + L_{\text{err}}$. Since the Fourier transform is linear, the same relationship holds for the Fourier representations of the signal components: $F\{L\} = F\{L_{\text{meas}}\} + F\{L_{\text{err}}\}$. Therefore, the high-frequency additions to the spectrum are given exactly by the Fourier transform of the saturation error. We call this component the *saturation noise*.

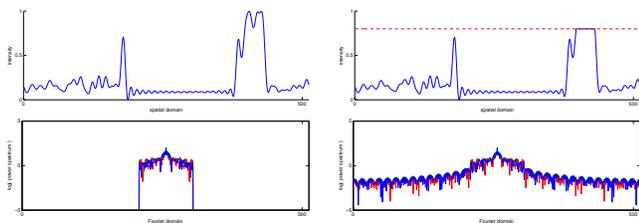


Figure 11: A band-limited signal (upper left), consisting of a single scanline taken from a high dynamic range image, and the same signal clipped at 0.8 maximum intensity level (upper right). The Fourier transform of the original scanline is band-limited (lower left), while the Fourier transform of the clipped version of the same scanline is heavily corrupted by high frequency components.

In the context of Fourier multiplexing, the higher frequency bands are also filled with information, i.e. with the different copies of the signal. The high-frequency noise introduced by saturation corrupts all copies encoded in the different Fourier tiles, so none of them can be reconstructed through straightforward tile extraction followed by the inverse transform.

Even more severe artifacts are introduced by full saturation of large image regions. Since the spatial structure of the super-pixels is removed completely in such regions, the local information cannot be copied into the high frequency bands of the image. Instead, the *DC* component of the signal, the constant part of the image, is increased. As a result, the copies, when transformed back to image space, exhibit black holes in image regions that would normally be saturated. An example artifact from one of the higher-frequency tiles is shown in Figure 12. It is worth to noting that this failure mode is a fundamental limitation of *all* Fourier multiplexing imaging techniques.

Usually, the masks employed by Fourier multiplexing techniques exhibit variations in transmissivity to generate repeating image space patterns. This implies that in general, there will also be image regions where only some parts of a super-pixel saturate, while others do not. We refer to these regions as being *partially* saturated. The effects of partial saturation are different from the fully saturated case. This can easily be seen by considering a partially saturated pattern as shown in Figure 13. The Fourier transform of a repeating pattern is a series of Dirac peaks whose magnitudes are determined by the Fourier transform of this pattern [Lanman et al. 2008]. By changing this pattern, partial saturation has the effect of changing the magnitudes of the Dirac peaks employed for Fourier

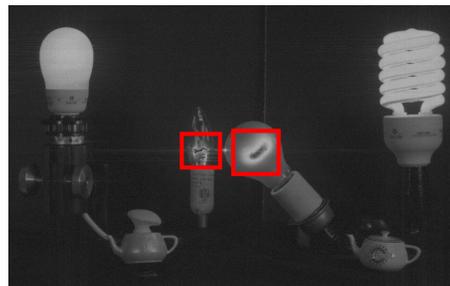


Figure 12: The impact of saturation noise in an image reconstructed from a single tile in Fourier space. The multiplexed image was taken with an array neutral density filters.

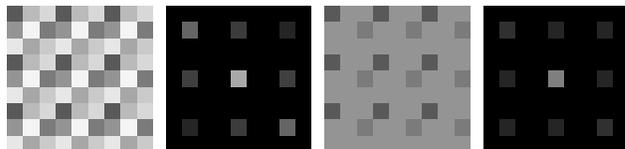


Figure 13: A 3×3 pattern of neutral density filters and its Fourier transform (left pair), partial saturation of this pattern causes changes in the magnitude of the Dirac peaks in the mask (right pair).

multiplexing. Since only parts of the images are affected by this type of saturation, the effect is not directly quantifiable. However, this example shows that the overall scale of copies obtained from Fourier multiplexing can have changing magnitudes as soon as sensor saturation is involved.

In summary, saturation has a significant impact on the performance of Fourier multiplexing techniques. Saturation, if not dealt with properly, can introduce severe artifacts in multiplexed information. However, in the next section we demonstrate that the additional Fourier components added into the signal can be considered data for a dynamic range reconstruction algorithm.

7 Fourier-Space Dynamic Range Optimization

Based on our previous analysis, we now introduce a Fourier-space optimization approach to recover dynamic range information in saturated image regions. The optimization is based on the idea that monochromatic neutral density filter masks can be used to create (differently scaled) copies in Fourier space. As we have seen, these copies are corrupted even if only the most transmissive among the neutral density filters saturates.

Note however, that we have two pieces of information about the signal we are recording. First, the original signal before modulation is band-limited. Second, the filter mask creates Fourier tiles with different (known) amplitudes, corresponding to the coefficients of the spatial Fourier basis. We are thus able to formulate an error measure ϵ in Fourier space that incorporates these two constraints:

$$\epsilon = \sum_{i=1}^{N-1} \sum_{j=i+1}^N \|s_i T_i(f_x, f_y) - s_j T_j(f_x, f_y)\|_2^2, \quad (8)$$

where $T_i(f_x, f_y) = F_i\{(1/s_i) L_s(x, y) + \eta_i\}$ is a tile describing a single copy of the signal in the Fourier domain, and η_i is sensor

noise. F_i is the Fourier transform that maps a full-resolution image from the spatial domain into the subset of the frequency space that is spanned by tile T_i . The scaling factors s_i describe the relative amplitudes of the individual tiles. For simplicity of notation we assume in the following that all tiles have been normalized by dividing through the corresponding factor.

We are now splitting the desired, mask-modulated target image L_s into a clipped part where the sensor does not saturate, and a second part where it does. That is, $L_s = L_{\text{unsat}} + L_{\text{sat}}$, where

$$L_{\text{unsat}}(x, y) = \begin{cases} L_s(x, y) & ; L_s(x, y) < L_{\text{max}} \\ 0 & ; \text{else} \end{cases}$$

$$L_{\text{sat}}(x, y) = \begin{cases} 0 & ; L_s(x, y) < L_{\text{max}} \\ L_s(x, y) & ; \text{else} \end{cases} .$$

The corresponding relationship in Fourier space is: $F\{L_s\} = F\{L_{\text{unsat}}\} + F\{L_{\text{sat}}\}$. The individual tiles are now given as

$$T_i = F_i\{L_{\text{unsat}}\} + F_i\{L_{\text{sat}}\} + \eta_i. \quad (9)$$

The term $F_i\{L_{\text{unsat}}\}$ can readily be computed from the captured image and represents measured data. $F_i\{L_{\text{sat}}\}$ includes the unknown variables (the non-zero subset of L_{sat}) that, in the presence of saturation, will cause the saturation noise in the Fourier domain. Combining Equations 8 and 9 yields

$$\varepsilon = \sum_{i=1}^{N-1} \sum_{j=i+1}^N \|s_i F_i\{L_{\text{unsat}}\} - s_j F_j\{L_{\text{unsat}}\} + s_i F_i\{L_{\text{sat}}\} - s_j F_j\{L_{\text{sat}}\} + \eta_i + \eta_j\|_2^2, \quad (10)$$

We assume that the sensor noise η_i is independently distributed in the spatial domain and observes a Gaussian noise distribution in the per-pixel image intensities. Thus, $F\{\eta_i\}$ has a uniform power spectrum with a Gaussian characteristic in each Fourier coefficient. This noise model, which is consistent with photon shot noise, allows us to use a quadratic error norm for optimization in Fourier space.

We encode ε in a linear system of equations, where we optimize the spatial pixel intensities of L_{sat} using an error measure defined in Fourier space. We show a simplified example assuming one copy at the DC peak and one copy of equal scale in a higher frequency band:

$$\min \left\| \begin{bmatrix} 1 & -1 & 0 \\ 0 & 0 & 0 \\ 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} F_1 \\ F_{DC} \\ F_1^* \end{bmatrix} (L_{\text{unsat}} + L_{\text{sat}}) \right\|_2^2. \quad (11)$$

We then formulate Equation 11 in matrix notation as $\min \|RF(L_{\text{unsat}} + L_{\text{sat}})\|_2^2$ with R encoding the relations between the different Fourier copies, and F performing the transformation from the spatial domain into the single frequency tiles. Since we are minimizing differences, matrix R is not of full rank. We deal with this problem by adding a regularizer S on the spatial components of the combined signal. In the regularizer we encode a spatial smoothness constraint, i.e. a curvature minimizing term, which is a reasonable assumption for the band-limited signal we are aiming to recover:

$$\min \|(RF + \alpha S)(L_{\text{unsat}} + L_{\text{sat}})\|_2^2. \quad (12)$$

Differentiating Equation 12 with respect to L_{sat} and setting the gradient to zero, we obtain a least squares description of our error measure:

$$(F^* R^* RF + \alpha S^* S) L_{\text{sat}} = - (F^* R^* RF + \alpha S^* S) L_{\text{unsat}}. \quad (13)$$

Note that the right hand side of this system is constant and represents our image space measurements.

We solve this linear system using the conjugate gradient for least-squares (CGLS) algorithm [Hansen 1998]. CGLS is a variation of the conjugate gradient method that solves the normal equations $A^* Ax = A^* b$ of a linear equation system $Ax = b$ without actually forming the matrix $A^* A$. Furthermore, matrix multiplications are handled transparently and can be replaced by fast image processing routines instead of constructing the actual matrices. Doing so facilitates the processing of full resolution images.

Optimization Results. An example scene containing saturated regions is shown in Figure 14. The upper left part shows magnifications of the multiplexed image captured with our scan camera prototype (2628×1671 pixels). Three of the Fourier tiles are seen above the Fourier transform of the image. Note that they are differently affected by saturation and sensor noise, as well as scanner artifacts. Figure 14 (center) shows the tone-mapped result of our HDR reconstruction. In order to compute this result, we performed our optimization on a grid of 3×3 tiles in the Fourier domain, each with a resolution of 876×557 pixels. The dynamic range of the captured scene is extended by a factor of 1.58 in this case. Note that this factor is with respect to the already increased dynamic range obtained by using the same neutral density filters array in combination with spatial reconstruction [Nayar and Mitsunaga 2000].

For validation, we compare our result with a high dynamic range ground truth image that was generated by combining 12 exposures of an SLR camera located next to the scan camera. As seen in the multi-exposure sequence on the right side of Figure 14 (center row), the dynamic range can be faithfully recovered. Note in particular the structure recovered in the cold fluorescent light bulb. The depicted LDR image (top row) was photographed using our scan camera without the attenuation mask and lacks details in bright image regions. The SLR image is shown at the bottom.

Another example scene is presented in Figure 15. Here, we also show results of our optimization for saturation of dark regions (highlighted in green) in addition to saturation in bright image parts (red highlights). As depicted by the color coded magnifications, the mask pattern is almost completely canceled out in regions that exhibit full saturation. Our optimization algorithm can restore the mask modulation in saturated bright parts as well as in dark parts.

Due to a slight mis-registration (rotation and shift) of the filter in front of the sensor in our prototype camera, as well as dust and scanner sensor artifacts, the point spread functions (PSFs) of the filter tiles in the Fourier domain do not exactly correspond to the filter specification before the print. In order to calibrate for these effects, we estimate the PSFs of the individual filter tiles in the Fourier domain of a calibration image.

Comparison to Previous Work. Figure 16 shows a comparison of our Fourier space optimization, and a spatial reconstruction

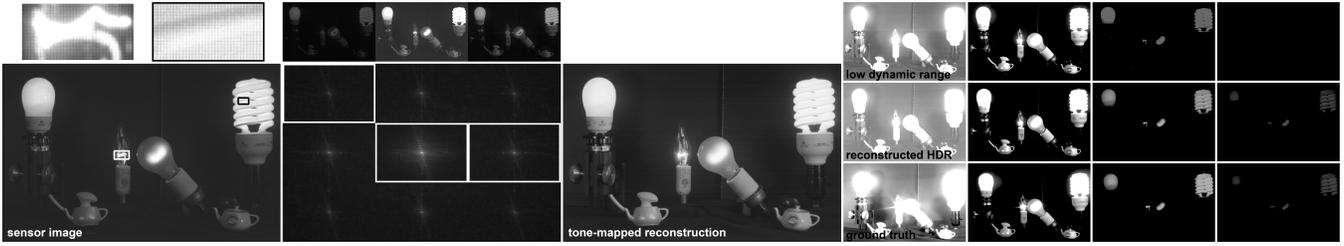


Figure 14: Mask modulated LDR image captured with our scan camera (left). The pattern introduced by the mask and saturated regions are enlarged on the upper left. The Fourier transform of the captured image contains several copies, which can individually be transformed into the spatial domain and reveal saturation artifacts. The right columns show several exposures of an unmodulated LDR image (top row) and our reconstruction (center) compared to ground truth images taken with an SLR camera (bottom).

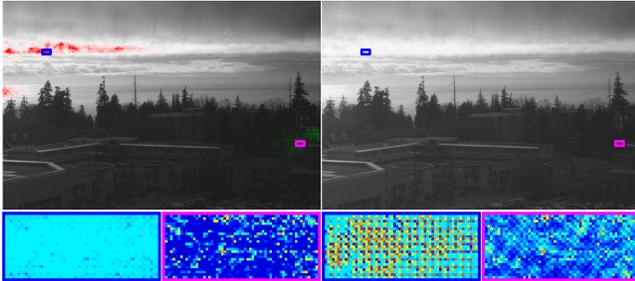


Figure 15: Outdoor scene captured with our scan camera prototype. Left: sensor image with saturated parts in red (bright) and green (dark). Right: reconstructed HDR image. The magnified parts show that our optimization recovers the mask pattern in completely saturated bright and dark image regions.

along the lines of the work by Nayar and Mitsunaga [2000]. For this experiment we simulated a one-dimensional sensor and used the same mask for both reconstruction methods. The dynamic range recovered with cubic interpolation as employed by the spatial Assorted Pixels approach corresponds to the the dynamic range of the image sensor times the dynamic range of the neutral density filter array.

As before, the Fourier reconstruction uses the fact that the spatially varying pixel mask implicitly creates copies in the Fourier domain. Although the Fourier-based approach initially suffers from the saturation artifacts discussed in Section 6, our optimization procedure manages to recover detail in regions not recovered by the Assorted Pixels approach.

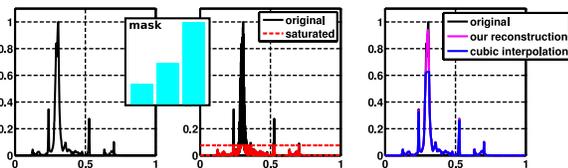


Figure 16: A band-limited 1D signal (left) is modulated with an attenuation mask pattern (one of the repeating tiles shown in inset), and captured by a sensor with a limited dynamic range (red dotted line, center). Our reconstruction (right, magenta) performs better than previously applied interpolation methods (right, blue).

8 Discussion and Conclusions

In this paper, we have derived a new framework for analyzing image multiplexing techniques. We show in particular a close relationship between spatially- and Fourier-multiplexed patterns. General spatial patterns such as the Bayer pattern can be reconstructed with a Fourier-space algorithm, while patterns derived for Fourier space can also be reconstructed by spatial interpolation, followed by a local de-correlation matrix of the individual channels. For the purposes of this paper, we have focused on the principal similarities of the two types of reconstruction. In practice, issues such as noise amplification, numerical stability, and quality of spatial interpolation play a role in choosing one reconstruction method over the other. We leave an investigation of these aspects to future work.

Our derivation shows that we can apply both spatial and Fourier multiplexed encodings for all dimensions of the plenoptic function. Our prototype implements multi-spectral imaging, but other dimensions such as time are also possible. We plan to investigate spatio-temporal multiplexing in the future, which could lead to new high-speed camera modes. Another topic of investigation would be combinations of multiple plenoptic dimensions.

Finally, we have demonstrated that conventional Fourier reconstructions of multiplexed images suffer from artifacts if the image sensor saturates. We derived a noise model to analyze these artifacts, and devised a Fourier-space optimization strategy for recovering dynamic range in clipped regions. Doing so, we can recover images with higher dynamic range than corresponding spatial reconstruction methods.

However, it should be noted that this optimization strategy requires known relationships between the Fourier tiles of the multiplexed image, in our case the property that the tiles are scaled versions of each other. This property does not hold for general plenoptic multiplexed images such as encoded light fields. We believe that investigating optimization algorithms for such encodings will be an interesting avenue for future work.

9 Acknowledgments

This work was supported by Dolby under the Dolby Research Chair in Computer Science at the University of British Columbia. Ivo Ihrke was supported by a Feodor-Lynen Fellowship of the Humboldt-Foundation, Germany.

We especially would like to thank Veeraraghavan et al. [2007] and Levoy et al. [2006] for making their data sets publicly available.

A Spectral Basis Decoding

When capturing a multiplexed image with a mask encoding $\{b_{k,l}^c, b_{k,l}^s\}$ as shown in Equation 3, the Fourier transform of the multiplexed image L_s contains tiles that represent filtered versions of the original radiance. However, due to the symmetries of the Fourier transform, the individual tiles are not completely independent in their encoding. Specifically, tiles $T_{k,l}$ and $T_{-k,-l}$ together encode the radiance image $L_{k,l}^c$ filtered through the spectral basis functions $b_{k,l}^c$, and $L_{k,l}^s$ filtered through $b_{k,l}^s$. Specifically, we have

$$L_{k,l}^c = F^{-1} \left\{ \frac{1}{2} (T_{k,l} + T_{-k,-l}) \right\}$$

and

$$L_{k,l}^s = F^{-1} \left\{ \frac{1}{2} i (T_{k,l} - T_{-k,-l}) \right\}.$$

Tile $T_{0,0}$ describes the DC component, the Dirac peak of which is always real. Therefore, it encodes only a single copy of the signal.

References

- ADELSON, E. H., AND WANG, J. Y. A. 1992. Single Lens Stereo with a Plenoptic Camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 2, 99–106.
- AGGARWAL, M., AND AHUJA, N. 2004. Split Aperture Imaging for High Dynamic Range. *Int. Journal of Computer Vision* 58, 1, 7–17.
- ALLEYSON, D., SÜSSTRUNK, S., AND HÉRAULT, J. 2005. Linear Demosaicing inspired by the Human Visual System. *IEEE Trans. Im. Proc.* 14, 4, 439–449.
- BAYER, B. E., 1976. Color imaging array, July. US Patent 3,971,065.
- CHANG, E., CHEUNG, S., AND PAN, D. Y. 1999. Color filter array recovery using a threshold-based variable number of gradients. In *Proc. SPIE*, vol. 3650.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. ACM SIGGRAPH*, 369–378.
- DUBOIS, E. 2005. Frequency-Domain Methods for Demosaicking of Bayer-Sampled Color Images. *IEEE Signal Processing Letters* 12, 12, 847–850.
- GEORGIEV, T., INTWALA, C., BABACAN, S., AND LUMSDAINE, A. 2008. Unified frequency domain analysis of lightfield cameras. In *Proc. ECCV*, III: 224–237.
- GORTLER, S., GRZESZCZUK, R., SZELINSKI, R., AND COHEN, M. 1996. The Lumigraph. In *Proc. ACM SIGGRAPH*, 43–54.
- GREIVENKAMP, J. 1990. Color Dependant Optical Prefilter for the Suppression of Aliasing Artifacts. *Applied Optics* 29, 5, 676–684.
- HANSEN, P. C. 1998. *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.
- LANMAN, D., RASKAR, R., AGRAWAL, A., AND TAUBIN, G. 2008. Shield Fields: Modeling and Capturing 3D Occluders. *ACM Trans. Graph. (SIGGRAPH Asia)* 27, 5, 131:1–131:10.
- LEVIN, A., FERGUS, R., DURAND, F., AND FREEMAN, W. T. 2007. Image and Depth from a Conventional Camera with a Coded Aperture. *ACM Trans. Graph. (SIGGRAPH)* 26, 3, 70.
- LEVOY, M., AND HANRAHAN, P. 1996. Light Field Rendering. In *Proc. ACM SIGGRAPH*, 31–42.
- LEVOY, M., NG, R., ADAMS, A., FOOTER, M., AND HOROWITZ, M. 2006. Light field microscopy. *ACM Trans. Graph.* 25, 3, 924–934.
- LIPPMANN, G. 1908. La Photographie Intégrale. *Academie des Sciences* 146, 446–451.
- MCGUIRE, M., MATUSIK, W., PFISTER, H., CHEN, B., HUGHES, J. F., AND NAYAR, S. K. 2007. Optical Splitting Trees for High-Precision Monocular Imaging. *IEEE Comput. Graph. Appl.* 27, 2, 32–42.
- MITSUNAGA, T., AND NAYAR, S. K. 1999. Radiometric Self Calibration. In *Proc. Computer Vision and Pattern Recognition*, 374–380.
- NARASIMHAN, S., AND NAYAR, S. 2005. Enhancing Resolution along Multiple Imaging Dimensions using Assorted Pixels. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27, 4 (Apr), 518–530.
- NAYAR, S., AND MITSUNAGA, T. 2000. High Dynamic Range Imaging: Spatially Varying Pixel Exposures. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, 472–479.
- NAYAR, S., AND NARASIMHAN, S. 2002. Assorted Pixels: Multi-Sampled Imaging With Structural Models. In *European Conference on Computer Vision (ECCV)*, vol. IV, 636–652.
- NG, R. 2005. Fourier Slice Photography. *ACM Trans. Graph. (SIGGRAPH)* 24, 3, 735–744.
- RASKAR, R., AGRAWAL, A., AND TUMBLIN, J. 2006. Coded Exposure Photography: Motion Deblurring using Fluttered Shutter. *ACM Trans. Graph. (SIGGRAPH)* 25, 3, 795–804.
- RASKAR, R., AGRAWAL, A., WILSON, C. A., AND VEERARAGHAVAN, A. 2008. Glare Aware Photography: 4D Ray Sampling for Reducing Glare Effects of Camera Lenses. *ACM Trans. Graph. (SIGGRAPH)* 27, 3, 1–10.
- SCHECHNER, Y., AND NAYAR, S. 2003. Generalized Mosaicing: High Dynamic Range in a Wide Field of View. *International Journal on Computer Vision* 53, 3, 245–267.
- TALVALA, E.-V., ADAMS, A., HOROWITZ, M., AND LEVOY, M. 2007. Veiling Glare in High Dynamic Range Imaging. *ACM Trans. Graph. (SIGGRAPH)* 26, 3, 37:1–37:9.
- TUMBLIN, J., 2007. Personal communication.
- VEERARAGHAVAN, A., RASKAR, R., AGRAWAL, A., MOHAN, A., AND TUMBLIN, J. 2007. Dappled Photography: Mask Enhanced Cameras for Heterodyned Light Fields and Coded Aperture Refocussing. *ACM Trans. Graph. (SIGGRAPH)* 26, 3, 69.
- VEERARAGHAVAN, A., RASKAR, R., AGRAWAL, A., CHELLAPPA, R., MOHAN, A., AND TUMBLIN, J. 2008. Non-Refractive Modulators for Encoding and Capturing Scene Appearance and Depth. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- WANG, S., AND HEIDRICH, W. 2004. The Design of an Inexpensive Very High Resolution Scan Camera System. *Computer Graphics Forum (Eurographics)* 23, 10, 441–450.