A New Reconstruction Filter for Undersampled Light Fields

J. Stewart,[1] J. Yu,[2] S.J. Gortler,[3] and L. McMillan [1]

[1] Department of Computer Science, The University of North Carolina at Chapel Hill
[2] Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology
[3] Division of Engineering and Applied Sciences, Harvard University

Abstract
This paper builds on previous research in the light field area of image-based rendering. We present a new reconstruction filter that significantly reduces the "ghosting" artifacts seen in undersampled light fields, while preserving important high-fidelity features such as sharp object boundaries and view-dependent reflectance. By improving the rendering quality achievable from undersampled light fields, our method allows acceptable images to be generated from smaller image sets. We present both frequency and spatial domain justifications for our techniques. We also present a practical framework for implementing the reconstruction filter in multiple rendering passes.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation — *Viewing algorithms*; I.3.6 [Computer Graphics]: Methodologies and Techniques — *Graphics data structures and data types*; I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture — *Sampling*

Keywords: Image-based rendering, light field, lumigraph, sampling, reconstruction, aliasing, multipass rendering

1. Introduction

In recent years, light field rendering has been offered as an alternative to conventional three-dimensional computer graphics. Instead of representing scenes via geometric models, light fields use a collection of reference images as their primary scene representation. Novel views can then be reconstructed from these reference images.

Conceptually, light fields are composed of a ray database, or more specifically, a database of radiance measurements along rays. In practice, these radiance measurements are usually organized as a set of camera images acquired along the surface of a parameterized two-dimensional manifold, most often a plane[1,2]. This leads to a four-dimensional description for each sampled ray (typically two for specifying the manifold coordinates of the camera, and two for specifying the image coordinates of each ray).

Since representing all light rays present in a scene is usually impractical or impossible, the database contains only a finite sampling of the rays. Thus, as with any discrete sampling of a continuous signal, we are faced with attempting to avoid aliasing through proper reconstruction. In general, aliasing can be introduced during initial sampling and during reconstruction. In this paper, we focus on aliasing introduced due to insufficient initial sampling, specifically undersampling along the two camera-spacing dimensions. Undersampling of the camera plane is common (and in some sense desirable) because the "samples" are actually high-resolution images requiring significant memory. Thus, a sparse sampling of the camera plane saves disk storage and reduces run-time memory requirements.

Of course, the problem with undersampling the camera plane is that aliasing is introduced. If an undersampled light field is rendered using the common linear interpolation method (referred to in this paper as "quadrilinear reconstruction"), then the results will exhibit an aliasing artifact called "ghosting", where multiple copies of a single feature appear in the final image. In an animation, these artifacts are not coherent from frame to frame, causing a distracting flicker. Examples are shown in Plate 1a and in the supplementary video. It is this ghosting that we are primarily concerned with in this research.

The principal contribution of this paper is the description of a new reconstruction filter that significantly reduces ghosting artifacts while maintaining sharp reconstruction at a user-selectable depth. Our reconstruction approach employs simple linear filters, and does not significantly impact computation cost of light field construction and rendering. Previous reconstruction filters have had the property that they are direct analogs of a real-world camera model with a fixed resolution and aperture. Our reconstruction filter departs from this tradition, in that it is not realizable by any single optical system. This departure allows us to combine the best properties of multiple realizable optical systems.

2. Background and Previous Work

As with standard rendering methods, attempting to limit aliasing artifacts is a significant problem in light field rendering. Levoy and Hanrahan[1] showed how light field aliasing can be eliminated with proper prefiltering. This prefiltering is accomplished optically by using a non-pinhole camera with an aperture that is at least as large as the spac-

ing between cameras[3]. Alternatively, prefiltering can be accomplished computationally by initially oversampling along the camera-spacing dimensions, and then applying a discrete low-pass filter, which models a synthetic aperture. These prefiltering techniques are effective in reducing aliasing. However, Levoy and Hanrahan recognized that the sampling density of the camera plane must be relatively high to avoid excessive blurriness in the reconstructed images[1]. Moreover, prefiltering has the undesirable side effect of forcing an a priori decision as to what parts of the scene can be rendered in focus during reconstruction[4].

Aliasing problems can also be ameliorated with the introduction of approximate depth information as suggested by Gortler et al.[2] An analysis of the tradeoffs between sampling density and depth resolution was presented by Chai et al.[5] Both papers specifically address the issue of improving reconstruction at lower sampling densities. In this work, however, we assume that the acquisition of suitable depth information for real-world scenes is difficult (or at least very inconvenient). Thus, we seek to avoid the use of depth information in our reconstructions.

Recently, two alternate approaches for dealing with aliasing in light fields have been proposed. Chai et al.[5] suggested that a sufficient condition for avoiding aliasing artifacts altogether is to limit the disparity of all scene elements to ±1 pixel. One contribution of the paper was the recognition that this condition places inherent constraints on camera sampling density. If the sampling density is sufficiently high, then the focal plane can be placed at the "optimal depth" and the resulting quadrilinear reconstruction will not exhibit ghosting artifacts. This optimal constant depth is found by computing the average of the minimum and maximum disparity, given by the following equation[5]:

$$z_c = \frac{1}{disp_c} = \left[\frac{1}{2} (disp_{min} + disp_{max}) \right]^{-1} = \left[\frac{1}{2} \left(\frac{1}{z_{max}} + \frac{1}{z_{min}} \right) \right]^{-1} \quad (1)$$

For an undersampled light field, the disparity of some scene elements is greater than ±1 pixel, and thus a quadrilinear reconstruction will contain aliasing artifacts. This can be corrected by applying a decimation filter (a low-pass filter followed by subsampling) to each source image. The disparity is reduced to ±1 pixel in these lower resolution images. In this paper, we call this process "band-limited reconstruction." A major disadvantage of this approach is that it can result in excessive blurring when compared to the original source images.

Alternatively, the dynamic reparameterization approach of Isaksen et al.[4] demonstrated that it is possible to make any depth exhibit zero disparity, allowing any particular scene element to be reconstructed without ghosting artifacts. However, artifacts will be apparent for other scene elements whose disparity falls outside of the ±1 pixel range. This problem is addressed by increasing the spatial support, or aperture size, of the reconstruction filter. This has the effect of diffusing the energy contributions of scene elements with large disparities over the entire output image.

However, this "wide-aperture reconstruction" introduces two new problems. First, scene elements away from the focal plane are blurred. Second, view-dependent variations in reflectance are lost. This is unfortunate since the
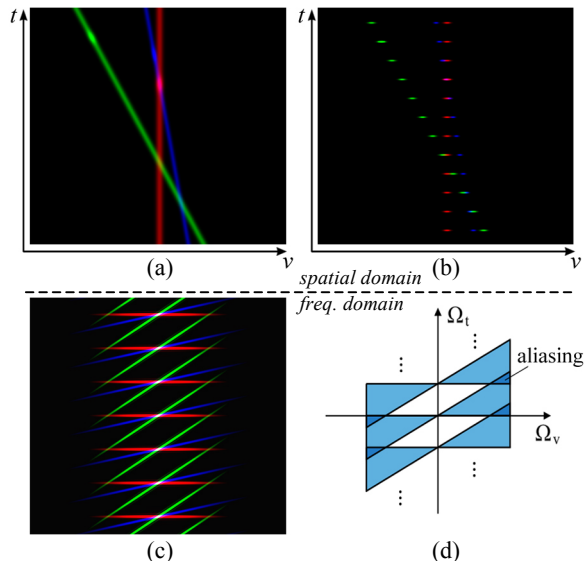
Figure 1: *Example undersampled light field.* (a) *Continuous 2D light field with three features at different depths.* (b) *The same light field sparsely sampled in the t dimension.* (c) *The resulting spectrum. Undersampling causes aliasing, indicated by the overlap of the red and green spectral lines.* (d) *The corresponding spectral diagram.*

accurate rendering of view-dependent effects, such as specular highlights, is a primary advantage of the light field representation over other image-based approaches.

In this paper, we present a new linear, spatially invariant reconstruction filter that combines the advantages of the band-limited reconstruction approach of Chai et al.[5] with the wide-aperture reconstruction approach of Isaksen et al.[4] The resulting filter reconstructs images in which ghosting artifacts are diminished. The images maintain much of the view-dependent information from the band-limited approach. Blurring is reduced using information from the sharply focused features of the wide-aperture reconstruction. Focal plane controls are used to specify a particular range of depths that appear in sharpest focus.

3. Frequency Domain

This section presents a frequency-domain description of the proposed reconstruction filter. We begin with an illustrative example of aliasing using a 2D light field. Throughout this discussion, we will use the nomenclature of Gortler et al.[2] That is, the camera plane uses parameters s and t, and the focal plane uses parameters u and v. Referring to Figure 1, we form an epipolar-plane image (EPI)[6] by fixing s and u. The resulting 2D light field represents a scene with three features at different depths, as shown in Figure 1a. The bright regions along each line model view-dependent reflections. Note that, in this simple example, we are ignoring the effects of occlusions. Note also that color is used simply to distinguish between the three features.[†]

[†] To view in color, refer to the EGSR 2003 CD-ROM or visit http://www.cs.unc.edu/~stewart/lf-recon-egsr2003/
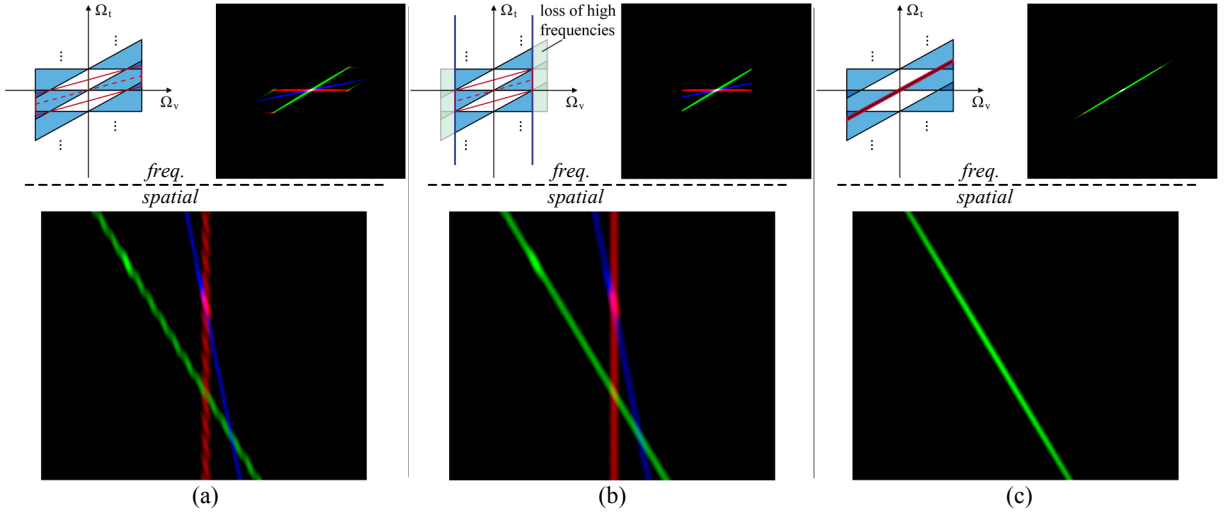
Figure 2: R*econstruction of undersampled light field.* (a) *Straightforward reconstruction. In the frequency domain, an ideal reconstruction filter is applied to the spectrum of Figure 1c. Note the red "tails" at the ends of the green spectral line. Note also the green at the ends of the red spectral line. This represents aliasing caused by undersampling. In the spatial domain, the blue feature lies near the optimal depth and is thus reconstructed well. However, the red and green features exhibit substantial artifacts.* (b) *Band-limited reconstruction after low-pass filtering. In the frequency domain, an ideal low-pass filter is applied to the aliased spectrum of Figure 1c. The skewed box filter from Figure 2a is then applied to the resulting band-limited spectrum. In the spatial domain, artifacts are reduced and some view dependence is maintained, but the image is blurrier than the original light field.* (c) *Wide-aperture reconstruction. In the frequency domain, the wide spatial-domain aperture results in a thin reconstruction filter. After applying the filter to the spectrum of Figure 1c, the spectral line for the green feature has been isolated. In the spatial domain, the green feature is reproduced with minimal ghosting artifacts and blurriness. Note, however, that the bright highlight representing view-dependent reflectance is spread along the line.*

Figure 1b represents a sparse sampling of the continuous light field in the *t* dimension. Figure 1c depicts the resulting power spectrum. Aliasing is evident due to the overlapping copies of the original signal's spectrum. Referring to Figure 2a, even an "ideal" box filter cannot properly recover the original light field, and thus ghosting artifacts appear in the reconstruction. We now examine various approaches for dealing with these artifacts.

3.1. Band-limited Reconstruction

One approach to reducing aliasing involves eliminating the spectral overlap via low-pass filtering. However, we must first calculate the filter width. If we perform quadrilinear reconstruction of an undersampled light field, and if we ignore view-dependent reflectance (i.e. if we assume a Lambertian BRDF model), then the spectrum is bounded by two lines, given in Chai et al.[5] by the following two equations:

$$\frac{f}{z_{min}}\Omega_v + \Omega_t = 0 \qquad \frac{f}{z_{max}}\Omega_v + \Omega_t = 0 \qquad (2)$$

From the same paper, we know that sampling along the camera dimension, *t*, with interval Δt results in spectral replicas spaced $2\pi/\Delta t$ apart. With this information, we can derive the appropriate width for the low-pass filter from simple line intersection calculations:

$$w = 2\pi \ / f \left(\frac{\Delta t}{z_{min}} - \frac{\Delta t}{z_{max}}\right) \ = 2\pi \ / (disp_{max} - disp_{min}) = 2\pi \ / \Delta disp \qquad (3)$$

Equation 3 correctly computes the filter width for Lambertian scenes. For scenes with view-dependent reflectance, however, the Lambertian assumption does not apply. Zhang and Chen[7] suggested that the spectrum of a non-Lambertian feature is "thicker" than that of its Lambertian counterpart. Specifically, if the features at depths z_{min} and z_{max} exhibit view-dependent reflectance, then the spectra for the two features will no longer be given by the perfect lines in Equation 2. Rather, if we assume that the resulting non-Lambertian BRDF model is band-limited, then the spectra will spread slightly (i.e. they will have a measurable thickness). Thus, the intersections of spectral copies (used to derive Equation 3) become small areas rather than points. Consequently, Equation 3 will slightly overestimate the filter width for non-Lambertian scenes, and so a small correction is required for such scenes. We assume that this scene-dependent adjustment is made.

Returning to our 2D example, band-limited reconstruction is illustrated in Figure 2b. After applying the low-pass filter, artifacts are effectively reduced in the resulting reconstruction. Much of the view-dependent information has also been maintained. However, the loss of high frequency information results in a final image that is blurrier than the original EPI shown in Figure 1a. More generally, all reconstructions will be noticeably blurrier than the original input images (the horizontal line segments that are visible in Figure 1b).

Figure 3: *Frequency-domain analysis of wide-aperture reconstruction. Focal plane positioned at the depth of the green feature.* (a) *A relatively small aperture results in a "tall" reconstruction filter. Significant aliased energy from the red and blue features falls within the filter support. The resulting reconstruction will thus contain ghosts for the red and blue features.* (b) *A very wide aperture results in a "thin" reconstruction filter. All aliased energy falls outside the filter support except where the red and green spectral lines cross at high frequencies.*
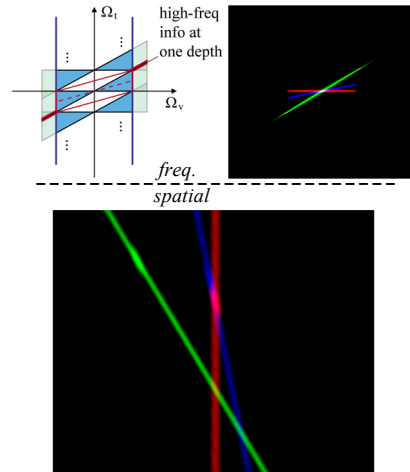

Figure 4: *Reconstruction using our method. In the frequency domain, the high-frequency information from the wide-aperture reconstruction of Figure 2c is added to the band-limited spectrum of Figure 2b. Note that the high frequencies have been recovered for the green feature. In the spatial domain, the green feature is reproduced in sharp focus and much of the original view dependence is still present. The blue and red features are blurry, but ghosting has been reduced.*

3.2. Wide-aperture Reconstruction

An alternative to band-limited reconstruction is the wide-aperture reconstruction approach of Isaksen et al.[4], which renders features at a single depth in sharp focus and reduces aliasing in the features at other depths by increasing the spatial support of the reconstruction filter. In contrast to quadrilinear reconstruction, which considers the contributions of only the nearest four cameras, the wide-aperture approach includes many cameras in the reconstruction (e.g. 64, 128, 256). This results in a synthetic aperture, with larger apertures produced by simply including more cameras.

Increasing the size of the synthetic aperture in the spatial domain decreases the "height" of the reconstruction filter in the frequency domain. Thus, for very wide apertures, the spectral information of a single feature can be extracted with minimal artifacts. This effect is demonstrated in Figure 3, which shows two possible filters applied to our example aliased spectrum of Figure 1c. In Figure 3a, the spatial support of the reconstruction filter is increased to include more than the four cameras used in quadrilinear reconstruction. The resulting frequency-domain filter is thinner than the corresponding quadrilinear filter, but it still includes aliased energy from the red and blue features.

In Figure 3b, the aperture size is increased further, producing a very thin filter in the frequency domain. In this example, nearly all of the aliasing energy falls outside the filter support except in the area where the red and green spectral lines cross at high frequencies. This remaining aliasing cannot be removed. However, for arbitrarily large apertures, the reconstruction filter can be made thin enough to reduce the aliased energy to imperceptible levels.

A wide aperture reduces the ghosting artifacts due to a single feature. It is, however, possible for a collection of periodically spaced features to conspire in such a way that traditional aliasing is exhibited in the reconstruction. The wide-aperture method depends on the diffusing effect of a single feature's reprojections onto random image regions. If, instead, these reprojections fall onto correlated regions, their combination might be reinforcing rather than diffusing. This potential limitation of the wide-aperture approach represents the light field equivalent of the proverbial "picket fence".

Another disadvantage of the wide-aperture approach is that view-dependent reflectance is greatly reduced. Recall from Section 3.1 that view dependence results in a "thickening" of the spectral lines. By reconstructing with a very thin filter, the wide-aperture method clips the spectrum such that its thickness approaches zero, resulting in a spectrum similar to that of a Lambertian feature.

Figure 2c illustrates the wide-aperture method applied to our example 2D light field. The focal plane is positioned to extract the desired feature, and a wide-aperture filter is applied, thereby isolating the spectral line for that feature. The thickness of this line is less than its corresponding value in the original spectrum. In the resulting spatial-domain reconstruction, the feature is reproduced with minimal ghosting artifacts and blur, but view-dependent reflectance is greatly reduced. Note also that the remaining two features have essentially vanished.

3.3. A New Reconstruction Approach

Our approach seeks to combine the advantages of wide-aperture reconstruction with those of band-limited reconstruction. The basic idea is to extract the high-frequency information present in the wide-aperture method and add it back into the result of the band-limited method. This combination enables a single feature to appear in sharp focus and maintain some view-dependent reflectance. The remaining features may be blurry, but they are free of ghosting artifacts and also maintain view-dependent reflectance.
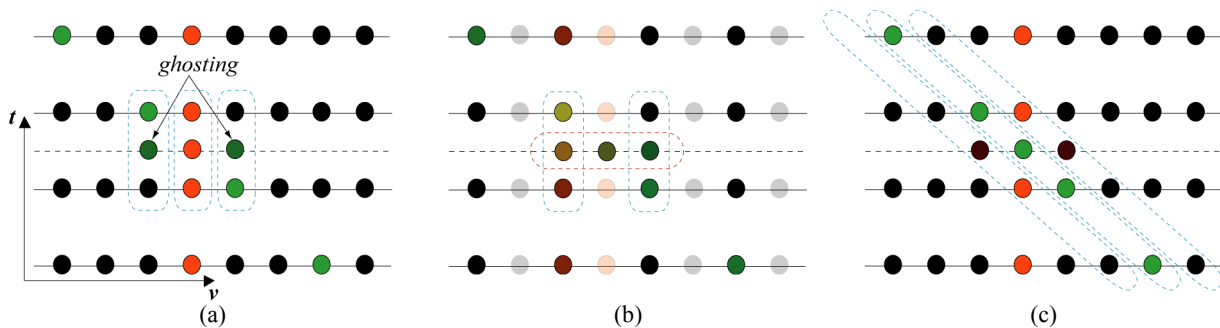
Figure 5: *Reconstruction comparison in the spatial domain. The t dimension represents camera position, while the v dimension corresponds to image pixels. The colored pixels roughly correspond to the red and green features of the example 2D light field. The horizontal dashed lines indicate the desired reconstruction camera position.* (a) *Quadrilinear reconstruction. Undersampling causes the disparity of the green feature to be greater than one pixel. Thus, the green feature appears twice in the reconstruction.* (b) *Band-limited reconstruction. The input images are low-pass filtered and downsampled, effectively reducing the disparity of the green feature to one pixel. However, low-pass filtering mixes the green feature with surrounding pixels, and the resulting reconstruction is thus a mix of the red feature, the green feature, and the black background.* (c) *Wide-aperture reconstruction. The focal plane is moved to the depth of the green feature. The result is a shearing of the light field such that the green feature becomes vertical, enabling it to be correctly reconstructed. Alternatively, one can view this as a shearing of the reconstruction filter such that it lines up with the green feature, as shown above.*

Our method is illustrated in Figure 4 for the example 2D light field. A high-pass filter (the complement of the low-pass filter from Figure 2b) is applied to the isolated spectral line from Figure 2c. The result represents the high-frequency information for the green feature. The addition of this spectral information to the band-limited spectrum of Figure 2b results in the spectrum shown in Figure 4. In the corresponding reconstruction, the recovered high-frequency information enables the green feature to be rendered in sharp focus, while maintaining the view-dependent highlight from the band-limited method.

In summary, our method can be described in the frequency domain as follows:

1. Calculate the width of the ideal low-pass filter using Equation 3
2. Apply the low-pass filter to the aliased spectrum
3. Apply the quadrilinear filter to the resulting band-limited spectrum
4. In a separate pass, apply the wide-aperture filter to the original aliased spectrum to isolate a particular feature
5. Apply a high-pass filter (the complement of the low-pass filter from Step 2) to the result
6. Add the high-frequency information from Step 5 to the result of Step 3

4. Spatial Domain

We now present a dual description of our method in the spatial domain. The intent is to provide a practical framework for implementing this new filter using multipass rendering and image processing operations. We begin by examining the problem.

Referring to Figure 5, if the light field is sufficiently sampled, then from scan line to scan line, a particular feature will shift a maximum of one pixel. Therefore, linear interpolation can correctly reconstruct all features without ghosting. Undersampling, however, causes some features to shift by more than one pixel between adjacent scan lines. In this case, linear interpolation generates duplications of these features, as shown in Figure 5a.

The band-limited approach (Figure 5b) addresses the problem by applying a low-pass filter to the input images, effectively reducing the disparity of all scene elements to ±1 pixel. Linear interpolation can then be used to generate reconstructions without ghosting. However, the resulting images will be blurry.

Alternatively, the wide-aperture approach (Figure 5c) conceptually applies a shear such that features at a particular depth have zero disparity. This allows sharp reconstruction of features at the chosen depth. Features not aligned with the shear are "aperture filtered" by combining samples from multiple cameras. Moving the focal plane changes the shear, and thus the depth that appears in sharpest focus can be dynamically selected at run time.

4.1. Implementation

To implement our method in the spatial domain, we first need to band-limit the light field via low-pass filtering. This is usually achieved by low-pass filtering the input images. However, since low-pass filtering and reconstruction via linear interpolation are both linear operators, they can be performed in either order. That is, one can blur the input images first and then perform reconstruction, or one can reconstruct output images first and then blur the result. Therefore, the first pass of our algorithm performs quadrilinear reconstruction with the focal plane at the optimal depth. The resulting image will contain ghosting artifacts, which are removed via low-pass filtering. This is analogous to the low-pass filtering in the frequency domain section.

The next rendering pass performs wide-aperture reconstruction as implemented by Isaksen et al.[4] The focal plane is positioned at run time to extract the desired feature. We then filter the result using the same low-pass filter as the previous pass.

Next, the unfiltered wide-aperture result is subtracted from the blurred version. This corresponds to the complement high-pass filter described in the frequency domain section. The resulting image contains the edge information for the selected depth. This edge information represents the high frequencies for that depth that were lost when we band-limited the light field in the previous pass.

The last step involves adding the edge image to the blurred band-limited reconstruction from the first pass, thereby restoring the high frequency information for a single depth. In summary, our method can be described in the spatial domain as follows:

1. Perform quadrilinear reconstruction
2. Low-pass filter the result
3. Perform wide-aperture reconstruction
4. Low-pass filter the result with the same filter used in Step 2
5. Subtract the unfiltered wide-aperture reconstruction from the filtered version to get high-frequency edge information
6. Add the edge image to the blurred band-limited reconstruction

Refer to Plate 1 for example images from various stages in the algorithm.

5. Results

The initial results presented in this section were achieved using the spatial-domain algorithm described in the previous section. The band-limited and wide-aperture reconstructions were produced with a ray-casting light field viewer. Separate code was then written to perform the image processing steps on these reconstructions.

We use a Gaussian kernel as our low-pass filter approximation, to avoid the ringing artifacts associated with truncated sinc filters (the spatial-domain equivalent of an ideal low-pass filter). Gaussians have many convenient properties that can be exploited in filtering including being radially symmetric, linearly separable, and having no non-negative weights. In addition, spatial-domain Gaussians conveniently map to Gaussians in the frequency domain. However, Gaussians do exhibit more blurring (passband attenuation) than other low-pass filter kernels. Alternatively, one could use a sinc filter in conjunction with standard windowing techniques, or one of many polynomial low-pass filter approximations[8].

Our first test scene consists of a texture-mapped quad in the foreground featuring the UNC Old Well. This image and the EGSR lettering are highly specular, simulating colored foil. A second quad, with a checkerboard texture, is located behind the foreground quad. Plate 1a represents quadrilinear reconstruction with the focal plane at the optimal depth. The camera plane of the test scene is highly undersampled and thus the resulting reconstruction exhibits extensive ghosting artifacts. This ghosting can be eliminated through low-pass filtering, but the resulting image is excessively blurry (Plate 1b).

Note that the near and far quad *could* be rendered via quadrilinear reconstruction without ghosting artifacts or excessive blurring, but the sampling density of the camera plane would have to be increased, resulting in more images

and run-time memory requirements. Instead of increasing the sampling density, our method combines the results of the band-limited reconstruction (Plate 1b) with high frequency information obtained from a wide-aperture reconstruction (Plate 1c through Plate 1e). The "wide aperture" encompasses 256 cameras, and thus each image contributes only 1/256 to the final reconstruction, allowing the aliased checkerboard, which is far from the focal plane, to be blurred below perceptible levels. Our result is shown in Plate 1f. Ghosting has been greatly reduced on the two quads, the foreground quad appears in focus, and the specular highlight on the Old Well has been preserved.

Note the faint "halo" around the foreground quad in Plate 1f. This artifact represents aliasing introduced during reconstruction. Specifically, since a Gaussian kernel only approximates an ideal low-pass filter, some high frequencies from adjacent spectral copies may still be present after filtering. A trade-off exists between decreasing this high-frequency leakage and minimizing unwanted passband attenuation. A wider spatial-domain kernel will reduce the leakage, but the resulting images will be blurrier. Refer to Table 1 for details on the kernels used in our results.

Our second test scene is an acquired light field. It contains objects over a wide range of depths: a bucket of pencils in the foreground; flowers, a stuffed animal, and a thermos in the mid-ground; and a curtain in the background. Again, undersampling causes ghosting in the quadrilinear reconstruction (Plate 2a), and this ghosting is removed through low-pass filtering in the band-limited reconstruction (Plate 2b). For the wide-aperture rendering (Plate 2c), the focal plane is positioned to extract the stuffed bear. The edge information from the sharp regions is then added into the band-limited image to produce our result (Plate 2d). We are able to maintain more detail in the foreground and background than the wide-aperture result, and we are sharper near the focal plane than the band-limited result. Note that a different choice of focal plane position during wide-aperture reconstruction produces a different result (Plate 2e and Plate 2f).

The effects of ghosting are sometimes hard to see in still images. In fact, one might prefer the ghosting in Plate 2a to the blurring in our results (Plate 2d and Plate 2f). However, in an animation, the ghosting will be incoherent from frame to frame, causing features to "jump around" in a distracting manner. Refer to the video that accompanies the paper for examples. The video also contains results for a third test scene.

6. Discussion

The light field reconstruction algorithm given in Section 4 is composed entirely of linear operations. Therefore, it is possible to achieve an identical filter, in principle, with a single pass using a fixed-weighted four-dimensional discrete convolution kernel. This filter would compute a weighted combination of rays within a hypervolume, and only a subset of these rays would have non-zero coefficients. Our two-pass approach is nearly optimal in that it considers only rays with non-zero weights, and with the possible exception of one ray, each ray is considered only once. It is, therefore, more efficient to implement our filter in two-passes as described. This is analogous to implementing a linearly separable kernel in two orthogonal passes.

An image generated by our light field reconstruction filter is not equivalent to any image that could be generated by a realizable camera. It could however, be simulated by combining the outputs of three cameras, where one image is focused on a target and captured at a high resolution with a large aperture. The second image would be focused at the optimal depth and captured at a lower resolution with a small aperture, while the third camera would capture a lower resolution with the same aperture and focus as the first camera. Even though our reconstruction filters are not analogous to any real-world optical system, they still exhibit desirable attributes. This begs the question of what other rendering effects might be achieved via linear reconstructions without physical analogs.

7. Conclusions

We have presented a new linear, spatially invariant reconstruction filter that reduces the ghosting artifacts of undersampled light fields. This is accomplished without increasing the camera-plane sampling density or requiring approximate scene geometry. We have also presented a practical framework for implementing these filters in the spatial domain, using two rendering passes.

Our approach combines the advantages of previous reconstruction methods. It provides more detail than either a band-limited reconstruction or a wide-aperture reconstruction. It also provides the flexibility of specifying what parts of the rendered scene are in focus.

By reducing the number of necessary camera-plane samples, this method allows for reconstruction with fewer images than previous techniques, thereby reducing storage space for the light field images and run-time memory requirements for the viewer.

References

1. M. Levoy and P. Hanrahan. Light Field Rendering. In *Proceedings of ACM SIGGRAPH 96*, 31-42, 1996.

2. S.J. Gortler, R. Grzeszczuk, R. Szeliski, and M.F. Cohen. The Lumigraph. In *Proceedings of ACM SIGGRAPH 96*, 43-54, 1996.

3. M.W. Halle. Holographic Stereograms as Discrete Imaging Systems. *Practical Holography VIII, SPIE 2176*: 73-84, 1994.

4. A. Isaksen, L. McMillan, and S.J. Gortler. Dynamically Reparameterized Light Fields. In *Proceedings of ACM SIGGRAPH 2000*, 297-306, 2000.

5. J. Chai, X. Tong, S. Chan, and H. Shum. Plenoptic Sampling. In *Proceedings of ACM SIGGRAPH 2000*, 307-318, 2000.

6. R.C. Bolles, H.H. Baker, and D.H Marimont. Epipolar-plane Image Analysis: An Approach to Determining Structure from Motion. *International Journal of Computer Vision*, 1(1): 7-56, 1987.

7. C. Zhang and T. Chen. *Generalized Plenoptic Sampling*. Electrical and Computer Engineering Technical Report AMP 01-06, Carnegie Mellon University, 2001.

8. D.P. Mitchell and A.N. Netravali. Reconstruction Filters in Computer Graphics. In *Computer Graphics (Proceedings of ACM SIGGRAPH 88)*, 22(4): 221-228, 1988.

	Size of light field	Resolution of source images	Resolution of output images	Disparity $\Delta disp$	Disparity of output images	Gauss std dev	Aperture size
Plate 1	16x16	1280x1280	1024x776	18 pixels	12 pixels	8 pixels	16x16
Plate 2	16x16	600x455	512x388	23 pixels	14 pixels	7 pixels	16x16

Table 1: *Parameters for the data sets used to test the algorithm. The $\Delta disp$ values are measured in the source images as the difference between the maximum and minimum disparity ($disp_{max}$ - $disp_{min}$). This provides an indication of the degree of undersampling (higher values indicate more severe undersampling). The column for "Disparity of output images" contains the maximum disparity for the rendered optimal-depth reconstructions. This value, $disp_{od}$, is used as a starting point in determining the standard deviation for the Gaussian filter kernel ($\sigma = \frac{1}{2} disp_{od}$). This initial standard deviation is then adjusted to balance the trade-off between high-frequency leakage and unwanted passband attenuation (i.e. the trade-off between aliasing and blurring).*

(a)　　　　　　　　　　(b)　　　　　　　　　　(c)

(d)　　　　　　　　　　(e)　　　　　　　　　　(f)
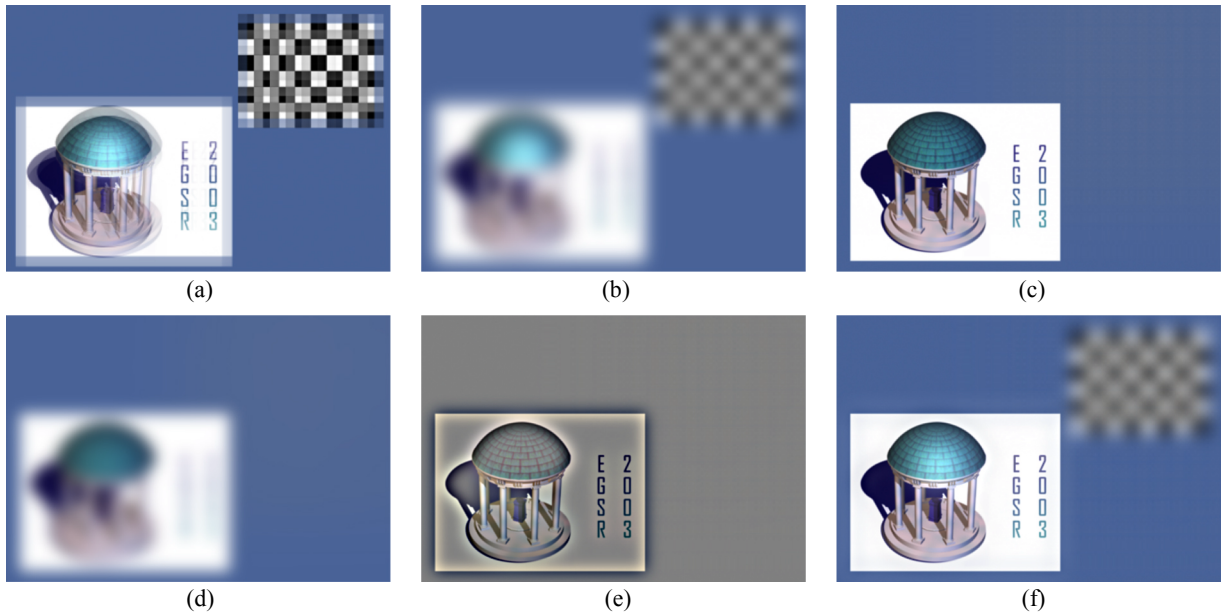
Plate 1: *Example results. (a) Quadrilinear reconstruction with the focal plane at the optimal depth. Ghosting results from undersampling. (b) Band-limited reconstruction via Gaussian blur of Plate 1a. Ghosting is reduced, and much of the highlight on the dome of the Old Well is maintained. (c) Wide-aperture reconstruction with 16x16 cameras. The foreground quad is in focus, but the highlight is gone. (d) Filtered wide-aperture reconstruction. (e) High frequencies of wide-aperture image (Plate 1d minus Plate 1c). (f) The final result (Plate 1b plus Plate 1e).*

(a)　　　　　　　　　　(b)　　　　　　　　　　(c)

(d)　　　　　　　　　　(e)　　　　　　　　　　(f)

Plate 2: *A comparison of the various reconstruction methods on an acquired light field. (a) Quadrilinear reconstruction with the focal plane at the optimal depth. Ghosting is visible around the bucket of pencils in the foreground, and in the backdrop curtain. (b) Band-limited reconstruction. Ghosting is reduced, but the result is blurry. (c) Wide-aperture reconstruction with the focal plane at the depth of the stuffed bear. Nearly all detail in the foreground pencils and the background curtain is lost. (d) Our results for this focal plane. Our method adds the high frequencies from the wide-aperture reconstruction to the band-limited result in Plate 2b. (e) Wide-aperture reconstruction with a focal plane at the depth of the foreground pencils. (f) Our results for this focal plane.*