

A 4-D Dual-Fan Filter Bank for Depth Filtering in Light Fields

Don Dansereau, *Member, IEEE*, and Len T. Bruton, *Fellow, IEEE*

Abstract—A light field is a four-dimensional (4-D) representation of the light permeating a scene—it parameterizes light rays as a function of position and direction. Such a structure can be measured using a specialized camera and can be used to render novel views of the scene it represents. It has previously been shown that the light field model of a scene may be filtered for a single depth by employing frequency-planar filters. Here, we show how a light field may be selectively filtered for a range of depths by forming a 4-D frequency passband that surrounds the intersection of two 4-D fans. A newly proposed cascaded filter bank is shown to approximate this passband, and its effectiveness is demonstrated on two scenes, both of which contain occlusions. Results are compared with those previously obtained using 4-D frequency-planar filters.

Index Terms—Filter banks, four-dimensional fan filters, four-dimensional signal processing, image analysis, light fields.

I. INTRODUCTION

IMAGE-BASED rendering has gained attention as a fast alternative to geometric model-based rendering. Light field rendering [1] and the lumigraph [2] are two similar image-based rendering techniques which seek to model a four-dimensional (4-D) subset of the more general seven-dimensional (7-D) plenoptic function [3] associated with a scene. In this way, the set of light rays permeating a scene is represented, rather than the geometry of the objects within the scene.

The 7-D plenoptic function describes the light rays in a scene as a function of position, orientation, spectral content, and time. This can be simplified to a 4-D function [1] by considering only the value of each ray as a function of its position and orientation in a static scene, and by constraining each ray to have the same value at every point along its direction of propagation. This disallows scenes in which the medium attenuates light as it propagates, and it fails to completely model the behavior of rays as they are occluded. These limitations are not an issue for scenes in a clear medium such as air, and for which the camera is not allowed to move behind occluding objects.

The 4-D light field typically parameterizes light rays using the two-plane parameterization (2PP), as depicted in Fig. 1. Each ray is described by its point of intersection with two reference




Fig. 1. Two-plane parameterization of light rays.

planes: the s - t plane given by $z = 0$, and the u - v plane, which is parallel to the s - t plane at some positive separation $z = d$. A full light field may consist of multiple sets of such planes, though this paper will deal only with a single set of reference planes. Note also that each sample of a light field can be taken as a grayscale intensity, though extension to utilize color samples given as red, green, and blue components is a simple matter of keeping one light field for each color channel, and repeating each operation accordingly, as is done in most image processing applications.

Because light fields accurately model scenes that are geometrically complex, they are well suited to act as an intermediary between the real world and computer vision algorithms. By storing a large amount of information about a scene prior to processing, light fields allow for simple algorithms to accomplish complex tasks.

Previous work on light field filtering has focused on filtering for a single depth in a scene [4] and on antialiasing filters [5]. This paper focuses on the task of extracting an arbitrary range of depths from a scene. Traditional approaches required dense estimation of a scene's depth [6], with computational complexity proportional to that of the scene's geometry. By using a light field, our approach remains independent of scene complexity. Furthermore, because of their simplicity, the techniques described are extremely robust and may have applications as diverse as robot navigation [7] and scene segmentation for object recognition, for example.

This paper deals exclusively with 4-D light fields; therefore, all signals are assumed to be four-dimensional. References to hyperplanes and planes refer to 4-D hyperplanes and planes in 4-D space, respectively. Also, wherever we refer to fan filters, we mean 4-D filters, for which the 4-D passband is fan shaped in two of the four dimensions, and independent of the remaining

Manuscript received October 25, 2004; revised March 20, 2006. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Trac D. Tran. The authors are pleased to acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Alberta Informatics Circle of Research Excellence (iCORE).

D. Dansereau is with SolVision Inc., Boucherville QC J4B 1E6, Canada (e-mail: dgdanser.engg@alumni.ucalgary.ca).

L. T. Bruton is with the Department of Electrical and Computer Engineering, University of Calgary, AB T2N 1N4, Canada (e-mail: bruton@ucalgary.ca).

Digital Object Identifier 10.1109/TSP.2006.885733

two dimensions. The continuous-domain light field will be denoted as $L_{\text{cont}}(s, t, u, v)$, and the discrete-domain version as $L(n_s, n_t, n_u, n_v)$, where \mathbf{n} is the discrete domain index of the signal. The continuous Fourier transform of the light field will be denoted as $L_{\text{freq}}(\Omega_s, \Omega_t, \Omega_u, \Omega_v)$. Note that the proposed filters all assume equal sample rates in all four dimensions, with differing sample rates requiring appropriate modifications.

The remainder of this paper is organized as follows: Section II explores the spectral characteristics of light fields and sets the stage for development of the 4-D filter in Section III. Section IV gives a design example and shows the results of filtering two different scenes for different ranges of depths. Results are compared with those obtained using frequency-planar filters. Conclusions are drawn in Section V.

II. SPECTRAL CHARACTERISTICS OF LAMBERTIAN SCENES

The depth filtering method proposed in this paper requires an understanding of the frequency content of a scene and how it relates to the depths of the objects in the scene. This section contains a development of previously established facts [8], [9], but examines them in a new light. It is hoped that by starting from geometric principles, this development will yield a more intuitive understanding of the phenomena at play.

A. Omnidirectional Point Light Source

Fig. 2 depicts a two-dimensional (2-D) slice along s and u of a subset of the rays emanating from an omnidirectional point source of light at the position $\mathbf{P} = [P_x, P_y, P_z]$. It is clear from this figure that for any given point on the s - t plane, there is only one point on the u - v plane for which a ray will intersect the light source. The result is that an s - u slice of the corresponding continuous-domain light field $L_{\text{cont}}(s, t, u, v)$ takes the form of a line, as depicted in Fig. 2(b). The equation of this line is given by

$$\left(\frac{d}{P_z} - 1\right)s + u = \frac{P_x d}{P_z} \quad (1)$$

where d is the separation of the reference planes. The behavior in the t and v dimensions is similar and can be given by

$$\left(\frac{d}{P_z} - 1\right)t + v = \frac{P_y d}{P_z}. \quad (2)$$

In 4-D space, (1) and (2) are the equations of two hyperplanes with normals in the directions $\mathbf{D}_1 = [d/P_z - 1, 0, 1, 0]$ and $\mathbf{D}_2 = [0, d/P_z - 1, 0, 1]$, respectively. The set of points in the light field that satisfy both (1) and (2) belongs to a plane defined by the intersection of these two hyperplanes. Points belonging to this plane of intersection correspond to rays emanating from the point light source and so take on the value of that light source, while all other points in the light field have a value of zero. Thus, an omnidirectional point light source is a plane of constant value in the light field $L_{\text{cont}}(s, t, u, v)$, where this plane is the solution of (1) and (2).

The normals of the hyperplanes (1) and (2) depend only on the perpendicular distance P_z of the light source from the s - t




Fig. 2. (a) Two-dimensional slice of a point source of light shown with the two reference planes; (b) two-dimensional slice of the corresponding light field $L_{\text{cont}}(s, t, u, v)$.

plane, implying that a set of omnidirectional point light sources at a single depth P_z exists as a set of parallel planes in the light field $L_{\text{cont}}(s, t, u, v)$, each with constant value [10]. The region of support (ROS) of the 4-D Fourier transform of such parallel planes can be shown to be a single frequency-domain plane through the origin (see the Appendix), given as the intersection of the two frequency-domain hyperplanes

$$\Omega_s + \left(1 - \frac{d}{P_z}\right) \cdot \Omega_u = 0 \quad (3)$$

$$\Omega_t + \left(1 - \frac{d}{P_z}\right) \cdot \Omega_v = 0. \quad (4)$$

Thus, the frequency-domain ROS of L_{freq} is a single plane through the origin, where the plane is the solution of (3) and (4). Similar results are shown in [9].

B. Lambertian Surface at a Single Depth

A Lambertian surface is one with ideal diffuse reflectance—that is, the luminance for any given point on such a surface is independent of viewing angle [11]. In this sense, a single point on a Lambertian surface behaves similarly to an omnidirectional point light source, and a planar Lambertian surface parallel with the reference planes at some depth $z = P_z$ will have a light field similar to that for the set of omnidirectional point light sources described above. Every point in this light field will belong to a constant-valued plane, and every plane in this light field will be parallel. The ROS of L_{freq} is again a plane through the origin with an orientation which depends only P_z . This result is shown in [9].

C. General Lambertian Surface

A Lambertian surface with general orientation—that is, one that has elements at depths varying from some minimum Z_{\min} to some maximum Z_{\max} —will, by superposition, have a frequency-domain ROS that is given by allowing the P_z parameter in (3) and (4) to vary continuously from Z_{\min} to Z_{\max} . The result is that the ROS has a fan shape in Ω_s, Ω_u , as defined by (5), and it also has a fan shape in Ω_t, Ω_v , as defined by (6). The overall ROS, which we will refer to as a dual fan, is defined by the intersection of these two fans [9]. Fig. 3 shows a typical




Fig. 3. Frequency-domain ROS of a light field of a surface which exists near the second reference plane.

ROS for an object that lies somewhere closer to the u - v reference plane than to the s - t reference plane.

$$\Omega_s + \left(1 - \frac{d}{P_z}\right) \cdot \Omega_u = 0, \quad Z_{\min} \leq P_z \leq Z_{\max} \quad (5)$$

$$\Omega_t + \left(1 - \frac{d}{P_z}\right) \cdot \Omega_v = 0, \quad Z_{\min} \leq P_z \leq Z_{\max} \quad (6)$$

It is worth noting that for surfaces that lie between half the reference plane separation and positive infinity ($d/2 < P_z < \infty$), the fan shapes in (5) and (6) will have central axes within 45° of the u and v axes, respectively. This observation will simplify the process of designing a fan-shaped passband in Section IV.

D. Occlusion and Specular Reflection

The light field properties derived above apply to scenes containing Lambertian surfaces and no occlusions. As a result, the filters developed below based on these characteristics will be most effective when applied to this type of scene. However, because most scenes with occlusions and specular reflections have most of their energy in diffuse reflection and nonoccluded surfaces, the filters operate properly in the presence of occlusion and specular reflection. Undesired effects associated with filtering this type of scene may include the vanishing of specular reflections from passband signals, the appearance of specular reflections from stopband signals, and the appearance of attenuated stopband signals in place of occluded passband signals.

III. 4-D DUAL-FAN FILTER BANK (4-D-DFFB)

Previous work has focused on extracting planar surfaces aligned with the reference planes [4]. In order to extract the more general class of objects occupying a range of depths, we propose and design a set of 4-D filter banks to extract the appropriately oriented frequency-domain dual-fan passband. Multidimensional filter banks provide a recently proven method for the formation of complex passband shapes [12].

In order to form the dual-fan 4-D passband, we propose to cascade two fan filter banks that satisfy (5) and (6), respectively. This is similar to the approach taken in [4], in which a 4-D planar passband was created by cascading two frequency-hyperplanar filters that satisfied (3) and (4), respectively. Each of the two fan filter banks approximates a 4-D fan-shaped passband by separating the input signal into sub-bands and applying hyperplanar filters with appropriately selected bandwidths and normals. This




Fig. 4. Fan filter bank $T(S_s, S_u)$ and $T(S_t, S_v)$.

approach is similar to that taken in [12] to form a 3-D conical passband. Fig. 4 shows one of the two cascaded fan filter banks—the input light field is separated into N_b subband signals, and each subband is filtered with a hyperplanar filter prior to being recombined into the final output signal.


It should be noted that all the frequency-planes that correspond to objects having stopband depths pass through the frequency-domain origin, and thus intersect the dual-fan passband of the 4-D-DFFB at the origin. As a result, the low-frequency components of the undesired stopband signals will be transmitted, though high-frequency components will be removed. Thus, stopband signals will never be entirely removed, but rather attenuated, particularly in the high-frequencies, resulting in a blurry stopband signal.

A. 4-D Frequency-Hyperplanar Filter

The fundamental component of the proposed filter banks is the hyperplanar filter. In [4], a hyperplanar filter with arbitrary orientation is designed by including an extra spatial variable in the 3-D planar filter discussed in [13]. A simplification of this 4-D filter is possible if the passband is oriented with selectivity in only two dimensions, allowing an appropriately chosen 2-D beam filter applied throughout the light field to yield the desired hyperplanar passband. This 2-D simplification yields a speed increase of a factor of about 1.7 over the 4-D version. The input–output difference equation of the 4-D IIR hyperplanar filter, which has selectivity in the Ω_s, Ω_u directions, is given by

$$y(\mathbf{n}) = \frac{1}{b_{00}} \left[\sum_{i=0}^1 \sum_{j=0}^1 L(n_s - i, n_t, n_u - j, n_v) - \sum_{i=0}^1 \sum_{j=0}^1 b_{ij} \cdot y(n_s - i, n_t, n_u - j, n_v) \right] \quad (7)$$

where the b_{ij} coefficients are found by iterating through the four sign configurations of $b_{ij} = 1 + [(-1)^i N_u + (-1)^j N_s]/B$, with N as the normal of the passband hyperplane, and with B as the -3 -dB bandwidth of the filter. Note that N_t and N_v are zero. Extension to create a hyperplanar filter with selectivity in the Ω_t - Ω_v directions is analogous and straightforward. The

Fig. 5. Approximated fan-shaped passband in Ω_s - Ω_u .

continuous-domain Laplace transform transfer function of the hyperplanar filter in (7) is given as

$$T(s_s, s_u) = \frac{1}{1 + \frac{(N_s s_s + N_u s_u)}{B}} \quad (8)$$

which, when applied throughout the 4-D light field, approximates a hyperplanar passband with its normal given by \mathbf{N} .

This filter is practical-BIBO stable, but only for values of \mathbf{N} that are non-negative along all four components. In order to filter outside the first hexadecimant, the direction of iteration is reversed for all dimensions for which the corresponding component of the normal is initially negative, then those components of the normal are changed to be positive [13].

B. 4-D Fan Filter Banks

Each of the two fan filter banks $T(s_s, s_u)$ and $T(s_t, s_v)$ in cascade approximate a fan shape defined by (5) or (6). In order to form these fan shapes, the input signal is first separated into N_b subbands using the one-dimensional (1-D) perfect-reconstruction (PR) filters described in [12]. Next, a hyperplanar filter with an appropriate normal and bandwidth is applied to each subband, and the bands are combined into a reconstructed signal. The resulting approximate-fan passband in Ω_s - Ω_u is depicted in Fig. 5. What follows is the design for the fan filter in the Ω_s - Ω_u dimensions, satisfying (5), with extension to the Ω_t - Ω_v dimensions being straightforward.

Because of our previous observation that the fan is likely to be oriented with its axis near Ω_u , the filter bank is designed to separate the input signal into N_b subbands along that dimension. Extension to deal with fans oriented near the Ω_s axis is a matter of separating subbands along that dimension instead.

A set of N_b finite-impulse-response (FIR) bandpass filters with real coefficients and order M is utilized to effect the sep-

aration into subbands. Each filter is a modulated sinc function truncated with a rectangular window

$$h_{BP}^k(u) = \begin{cases} w(u)h_{LB}(u) \cos\left(\frac{2\pi k u}{L}\right), & k = 0, N_b - 1 \\ 2w(u)h_{LB}(u) \cos\left(\frac{2\pi k u}{L}\right), & k = 1, 2, \dots, N_b - 2 \end{cases} \quad (9)$$

where $L = 2(N_b - 1)$, and $w(u)$ is a rectangular window of length M , with M chosen such that an integer number of lobes of the sinc function are contained within the ROS—setting $M = 2Lq + 1$, where q is the number of lobes, satisfies this criterion. Low values of q (as low as $q = 1$), give less selectivity but a less jagged approximation to the fan shape. The L -band filter h_{LB} is given by

$$h_{LB}(u) = \frac{1}{L} \text{sinc}\left(\frac{(u - \frac{M-1}{2})}{L}\right). \quad (10)$$

With the input signal split into bands along Ω_u , all that remains is to apply the hyperplanar filters with correct bandwidths and normals. The normal for all the filters is the same, set to match the orientation of the central axis of the fan, given as the mean of the orientations of the planes corresponding to Z_{\min} and Z_{\max} . The angle formed between the Ω_u axis and the fan's central axis is given by

$$\theta_c = \frac{\tan^{-1}\left(\frac{d}{Z_{\min}} - 1\right) + \tan^{-1}\left(\frac{d}{Z_{\max}} - 1\right)}{2} \quad (11)$$

and the normal, which aligns the hyperplanar filters with this axis, is

$$\mathbf{N}_{su} = \frac{[1, 0, -\tan(\theta_c), 0]}{\sqrt{1 + (\tan(\theta_c))^2}}. \quad (12)$$

The bandwidth of each hyperplanar filter depends on the difference of the slopes of the hyperplanes corresponding to Z_{\min} and Z_{\max} . There are many ways to calculate the bandwidth—one way is to approximate it as the distance, at the center of each subband along Ω_u , between the minimum- and maximum-depth hyperplanes, along Ω_s . This approach yields the expression

$$B_k = \frac{d(k + 0.5)}{2N_b} \left[\frac{1}{Z_{\min}} - \frac{1}{Z_{\max}} \right] + c, \quad k = 0, 1, \dots, N_b - 1. \quad (13)$$

Note that the bandwidth will increase proportionally with the distance of the subband from the Ω_s axis, thus creating the desired fan shape. Note also that the addition of a positive constant offset c to B_k will result in a passband that *surrounds* the fan in a bowtie shape, leaving room for some deviation in the passband from the ideal fan shape—such deviation might be caused by occlusion or specular reflection, for example.

C. 4-D-DFFB

In order to construct the 4-D-DFFB, two frequency-fan filters are arranged in a cascaded configuration—one satisfying (5), and one satisfying (6). The overall transfer function is given by

$$T(s_s, s_t, s_u, s_v) = T_1(s_s, s_u) \cdot T_2(s_t, s_v) \quad (14)$$

where T_1 and T_2 are the transfer functions of the two fan filter banks. The magnitude frequency response of (14) is unity only

where both filters have a magnitude frequency response of unity—that is, only where (5) and (6) are both satisfied. Thus, a filter that approximates the intersection of two fan-shaped passbands in four dimensions is achieved.

D. Zero-Phase Filtering

The filter described above yields good depth discrimination but with a long nonideal transient response in the output light field, the effects of which are darkening of the light field at the edges, and smearing of the images in the direction of iteration. This is especially significant for light fields that have low sample rates in s and t —a typical light field might have only tens of samples in these dimensions—and thus a method to reduce the effects of the transient response is desirable.

The proposed technique is to employ zero-phase filtering: that is, to refilter in a second pass the output signal with the direction of iteration reversed along each dimension, as in [4]. This is equivalent to flipping the normal of the frequency-hyperplanar passband of each frequency-hyperplanar filter in the second pass. This has no effect on the magnitude frequency response of the second-pass filters but enforces a zero-phase frequency response, leading to a much shorter transient response. By keeping extra output samples as they are smeared off the edge of the light field by the first-pass filter bank, then utilizing them as input when applying the second-pass filter bank, the darkening of the edges of the light field can also be reduced significantly. Furthermore, because a second-pass filter bank is being applied, the overall magnitude frequency response is squared, significantly increasing selectivity. The disadvantages of the zero-phase filtering technique are increased processing time and increased memory requirements.

IV. RESULTS

A. Design Example

Assume that a light field with equal sample rates in all four dimensions and a plane separation $d = 24$ cm is to have an object extracted that lies between the depths 24 and 30 cm. Choosing a four-band filter bank that uses a single-lobed sinc function ($N_b = 4, q = 1$), we find $L = 6$ and $M = 13$. Equation (10) is used to build the L -band filter h_{LB} of order M , then (9) is used to build the four corresponding bandpass filters.

Working first in the Ω_s – Ω_u directions, the bandpass filters are applied with selectivity in the u direction to obtain four subband signals. According to (11) and (12), the center angle θ_c is about -0.1 rad, and the appropriate normal is approximately $N_{su} = [1, 0, 0.1, 0]/1.005$. The bandwidth for each band, for zero bandwidth offset ($c = 0$), is given by (13) as $B_k = 0.025(k + 0.5) \approx 1/80, 3/80, 5/80, 7/80$. The resulting filter coefficients are summarized in Table I.

Once the subbands have been filtered by the hyperplanar filters, they are summed together according to Fig. 4 in the first filter bank, then again separated into subbands in the second filter bank by applying the bandpass filters $h_{BP}^k(v)$. Given identical sample rates in all dimensions, the center angle and the bandwidths of the appropriate hyperplanar filters are identical to those used for the first filter bank. The normal is approximately

TABLE I
SAMPLE HYPERPLANAR FILTER COEFFICIENTS

Coeff.	$k=0$	$k=1$	$k=2$	$k=3$
b_{00}	88.6	30.2	18.5	13.5
b_{01}	72.6	24.9	15.3	11.2
b_{10}	-70.6	-22.9	-13.3	-9.2
b_{11}	-86.6	-28.2	-16.5	-11.5

TABLE II
LIGHT FIELD PARAMETERS

Color Channels	3	s, t Size (cm)	21
s, t Samples	32	u, v Size (cm)	15
u, v Samples	256	Separation d (cm)	45




Fig. 6. Input light field: A simple scene.

$N_{tv} = [0, 1, 0, 0.1]/1.005$. The summation of the filtered subbands is the filter bank output.

By repeating the above process with the direction of iteration for each hyperplanar filter reversed, a zero-phase filter is formed.

B. Filtering a Real-World Scene

A 4-D-DFFB was applied to two a light field with geometric parameters as summarized in Table II. The light field was measured using a gantry system and contains occlusions. Fig. 6 shows the input light field as a slice through u and v for values of s and t near the center of the light field. The background is a poster of a supernova imaged by the Dominion Radio Astrophysical Observatory, British Columbia, Canada. The foreground is a beer coaster mounted at approximately 45° to the reference planes—a wooden dowel can be seen holding the coaster in place. The poster is at a depth of 66 cm, and the coaster occupies a range of depths from about 40 to 50 cm. The geometry of this light field was chosen to emphasize




Fig. 7. Output for (a) 4-D-DFFB; (b) zero-phase 4-D-DFFB; (c) 4-D-DFFB with offset; and (d) zero-phase 4-D-DFFB with offset.




Fig. 8. Zero-phase frequency-planar filter output: (a) $B = 0.3$ and (b) $B = 0.1$.

the coaster. Since the coaster is mounted at an angle to the reference planes, this will demonstrate passband performance over a range of depths.

To the authors' knowledge, there are no other methods for employing linear filters for the extraction of a range of depths from a light field. The most closely related work extracts a single depth from the scene, as in the zero-phase frequency-planar filter designed in [4]. For the sake of comparison, that filter was applied with bandwidths 0.3 and 0.1—note that a bandwidth of 0.3 corresponds to the best fit of the planar filter to the ideal dual-fan passband.

The results can be seen in Figs. 7 and 8. All the filters were successful at attenuating the background, though clearly the zero-phase filters displayed more selectivity and a shorter transient response, as evidenced by an increase in the blurring of the stopband signal, and a decrease in the darkening near the light field edges. The 4-D-DFFB caused significantly less distortion in the passband than the frequency-planar filters, particularly near the extremities of the coaster. The filters with the constant bandwidth offset improved the performance near the extremities of the coaster significantly, possibly due to the




Fig. 9. Input light field: A more complex scene.




Fig. 10. Results of applying: (a) zero-phase DFFB and (b) FP filter.

occlusions associated with these locations, though these filters provided less stopband attenuation. It can be seen in Fig. 8 that the frequency-planar filters caused less blurring of the poster, as in (a), or more blurring of the coaster, as in (b), and so the 4-D-DFFB represents a significant improvement over the frequency-planar filter for extracting objects occupying a range of depths.

C. Filtering a More Complex Scene

A zero-phase 4-D-DFFB was also applied to a light field of a virtual scene with much more complex geometry than that of the light field from the previous section. The scene can be seen in Fig. 9—it contains two trees at different depths and a field of grass stretching off to the horizon. The filter was set to extract a range of depths centered on the further (leftmost) tree, and corresponding roughly to two thirds of its width. The results were compared to those obtained using a best-fit FP filter (results are shown in Fig. 10). As seen in both the grass and the foliage, the DFFB is much more effective at extracting a range of depths, with the FP filter causing serious attenuation of high frequencies, particularly as distance from the center of the desired range of depths increases.

$$\begin{aligned}
L_{\text{freq}}(\Omega) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} L_{\text{cont}}(s, t, u, v) e^{-j(\Omega_s s + \Omega_t t + \Omega_u u + \Omega_v v)} dv du dt ds \\
&= \delta(\Omega_s) \delta(\Omega_t) \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} L(\Omega_s, \Omega_t, u, v) e^{-j(\Omega_u u + \Omega_v v)} dv du.
\end{aligned} \tag{15}$$

V. CONCLUSION

A 4-D recursive filter bank—the 4-D-DFFB—was proposed and designed to extract objects over a range of depths in a light field model of a scene by cascading two appropriately designed fan filter banks. Zero-phase dual-fan filtering was accomplished by applying second-pass filters with reversed directions of iteration.

Results demonstrated the isolation of elements within a range of depths, blurring and attenuating all undesired elements, and significantly outperforming frequency-planar filters for pass-band signals occupying a range of depths.

Passband performance was improved through the addition of a constant offset to sub-band bandwidths, at the cost of weaker stopband performance. The proposed zero-phase technique significantly reduced the length of the transient response of the DFFB, while increasing selectivity, at the cost of increased memory requirements and processing time.

Specific applications of the filters presented here might include depth-from-focus [10] for robot navigation, segmentation of three-dimensional scenes for biomedical applications, or extraction of foreground objects in human-computer interface applications. The possibility of applying simple and robust techniques to accomplish complex tasks is exciting, and ongoing research into areas such as object recognition [14], [15], and scene modeling [16] is promising.

Future work might involve optimization, for operation in real-time, through the use of systolic architectures or multi-rate filtering. A further exploration of the effects of occlusions and specular reflections on the spectral support of light fields may also prove interesting.

APPENDIX

The 4-D ROS of a light field of a planar scene can be found by first dealing with the specific case in which all objects in the scene are on the reference plane ($P_z = d$), and then generalizing to other plane depths. For objects on the reference plane, L_{cont} is invariant along s and t , so the 4-D Fourier transform is (see (15) at the top of the page).

The ROS of $L_{\text{freq}}(\Omega)$ is the plane defined by the intersection of the two 4-D delta functions.

Generalizing to planar scenes at other depths will yield light fields that are invariant along two different orthogonal vectors—this can be viewed as a rotated version of the scenario above:

$$L'_{\text{cont}}(s, t, u, v) = L_{\text{cont}}(R[s, t, u, v]^T) \tag{16}$$

$$L'_{\text{freq}}(\Omega) = L_{\text{freq}}(R\Omega^T) \tag{17}$$

under which the ROS will become a rotated version of the plane defined by the intersection of the two 4-D delta functions in (15).

REFERENCES

- [1] M. Levoy and P. Hanrahan, "Light field rendering," in *Proc. SIGGRAPH '96*, 1996, pp. 31–42.
- [2] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen, "The lumigraph," in *Proc. SIGGRAPH '96*, 1996, pp. 43–54.
- [3] E. H. Adelson and J. R. Bergen, "The plenoptic function and the elements of early vision," in *Computation Models of Visual Processing*. Cambridge, MA: MIT Press, 1991, pp. 3–20.
- [4] D. Dansereau and L. T. Bruton, "A 4-D frequency-planar IIR filter and its application to light field processing," in *Proc. Int. Symp. Circuits Systems (ISCAS)*, Bangkok, Thailand, May 2003, vol. IV, pp. 476–479.
- [5] J. X. Chai, X. Tong, S. C. Chan, and H. Y. Shum, "Plenoptic sampling," in *Proc. SIGGRAPH 2000*, pp. 307–318.
- [6] I. J. Cox, S. L. Hingorani, and S. B. Rao, "A maximum likelihood stereo algorithm," *Computer Vis. Image Understand.*, vol. 63, no. 3, pp. 542–567, May 1996.
- [7] Heigl, Denzler, and Niemann, "Combining computer graphics and computer vision for probabilistic visual robot navigation," *Enhanced and Synthetic Vision 2000*, vol. 4023, Proceedings of SPIE, pp. 226–235, Apr. 2000.
- [8] C. Zhang and T. Chen, "Spectral analysis for sampling image-based rendering data," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 13, no. 11, pp. 1038–1050, 2003.
- [9] S. C. Chan and H. Y. Shum, "A spectral analysis for light field rendering," in *Proc. 2000 Int. Conf. Image Processing*, 2000, vol. 2, pp. 25–28.
- [10] A. Isaksen, L. McMillan, and S. Gortler, "Dynamically reparameterized light fields," in *SIGGRAPH'00*, 2000, pp. 297–306.
- [11] D. Hearn and M. P. Baker, *Computer Graphics, C Version*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1997.
- [12] L. T. Bruton, "Selective filtering of spatio-temporal plane waves using 3-D cone filter banks," in *Proc. PACRIM'01*, Aug. 2001, pp. 67–70.
- [13] L. T. Bruton and N. R. Bartley, "Three-dimensional image processing using the concept of network resonance," *IEEE Trans. Circuits Syst.*, vol. CAS-32, no. 7, pp. 664–672, Jul. 1985.
- [14] B. Heigl, J. Denzler, and H. Niemann, "On the application of light field reconstruction for statistical object recognition," in *Proc. Eur. Signal Processing Conf. (EUSIPCO)*, 1998, pp. 1101–1105.
- [15] R. Gross, I. Matthews, and S. Baker, "Fisher light-fields for face recognition across pose and illumination," in *Proc. German Symp. Pattern Recognition (DAGM)*, Sep. 2002, pp. 481–489.
- [16] C. Vogelgsang, B. Heigl, G. Greiner, and H. Niemann, "Automatic image-based scene model acquisition and visualization," in *Proc. Workshop Vision, Modeling and Visualization*, Saarbrücken, Germany, 2000, pp. 189–198.



Donald Dansereau (M'05) received the B.Sc and M.Sc. degrees in electrical and computer engineering from the University of Calgary, Calgary, AB, Canada, in 2001 and 2004, respectively. His Master's research focused on novel ways of applying multidimensional filtering techniques to four-dimensional (4-D) light fields, with applications in computer vision.

He currently works as a Junior Vision Engineer at SolVision, Inc., Boucherville, QC, Canada. His main research interests are in computer vision and robotics.

Mr. Dansereau received a Governor General's Gold Medal for his thesis work in 2004.



Len T. Bruton (M'71–SM'80–F'81) is a Professor of electrical and computer engineering at the University of Calgary, Calgary, AB, Canada, where he carries out research in the fields of analog and digital signal processing with emphasis on multi-dimensional (MD) circuits and systems applied to the design and implementation of two-dimensional (2-D), three-dimensional (3-D), and four-dimensional (4-D) filtering techniques. He is especially interested in the VLSI real-time applications of 2-D and 3-D spatio-temporal filters for array processing in such fields as directional wireless, video sequence processing, seismic and audio signal processing. He is also the inventor of a number of widely used filtering methods, including the class of frequency-dependent negative resistance (FDNR) analog filters and the class of lossless discrete integrator (LDI) discrete-domain filters.

Dr. Bruton is a member of the Royal Society of Canada, a recipient of the 2002 Education Award of the IEEE Circuits and Systems Society (CAS-S), the 50th Jubilee Medal of the IEEE CAS-S, and the 1994 Outstanding Engineer Award of IEEE Canada (Region 7) "for leadership in engineering and the engineering profession and for research in the field of microelectronic digital and analog real-time filter design." In Canada, he received the 1992 Manning Principal Award for Innovation, and he is one of 162 scientists selected for inclusion in the textbook *Great Canadian Scientists* by B. Shell (Vancouver, BC, Canada: Polestar Book Publishers, 1997). He received the 1992 Alberta Science and Technology Award (ASTech) for Innovation in Science in recognition of his leadership in science, and in 1993, he received the Federal Government of Canada's 125th Anniversary of Canadian Confederation Medal in recognition of his significant contributions to compatriots, to community, and to Canada.