

Artifact reduction in lenticular multiscopic 3-D displays by means of anti-alias filtering

Janusz Konrad and Philippe Agniel

Department of Electrical and Computer Engineering, Boston University
8 Saint Mary's Street, Boston, MA 02215 [jkonrad,pagniel]@bu.edu

ABSTRACT

This paper addresses the issue of artifact visibility in automultiscopic 3-D lenticular displays. A straightforward extension of the two-view lenticular autostereoscopic principle to M views results in an M -fold loss of horizontal resolution due to the subsampling needed to properly multiplex the views. In order to circumvent the imbalance between the horizontal and vertical resolution, a tilt can be applied to the lenticules to orient them at a small angle to the vertical direction, as is done in the *SynthaGram*TM display from Stereographics Corp. In either case, to avoid aliasing the subsampling should be preceded by suitable lowpass pre-filtering. Although for purely vertical lenticules a sufficiently narrowband lowpass horizontal filtering suffices, the situation is more complicated for diagonal lenticules; the subsampling of each view is no more orthogonal, and more complex sampling models need to be considered. Based on multidimensional sampling theory, we have studied multiview sampling models based on lattices. These models approximate pixel positions on a lenticular automultiscopic display and lead to optimal anti-alias filters. In this paper, we report results for a separable approximation to non-separable 2-D anti-alias filters based on the assumption that the lenticule slant is small. We have carried out experiments on a variety of images, and different filter bandwidths. We have observed that the theoretically-optimal bandwidth is too restrictive; aliasing artifacts disappear, but some image details are lost as well. Somewhat wider bandwidths result in images with almost no aliasing and largely preserved detail. For subjectively-optimized filters, the improvements, although localized, are clear and enhance the 3-D viewing experience.

Keywords: Automultiscopic displays, lenticular displays, multidimensional sampling, anti-alias filtering

1. INTRODUCTION

Stereoscopic displays have been effectively used for visualization of 3-D information for many years. Whether in the form of anaglyph, polarized or shuttered glasses, or by means of lenticular or parallax barrier mechanisms, stereoscopic visualization quality has been steadily improving. One reason for this improvement has been a continuous evolution of the underlying display technology, while the other – a steady progress in the application of digital signal processing methods. For example, a significant improvement in color reproducibility in anaglyph stereoscopic viewing has been achieved by means of digital signal processing,¹ while an efficient attenuation of view crosstalk (ghosting) in active-eyewear (liquid-crystal shutters) stereoscopic visualization has been achieved by means of suitable digital image pre-processing.² Further improvements can be expected as we understand better deficiencies of 3-D displays and seek signal processing means of defeating them.

Although stereoscopic displays are effective in conveying 3-D information to a static viewer, serious difficulties arise when the viewer starts to move, e.g., turns his/her head. In these circumstances, motion-parallax conflict arises since the screen parallax does not change with viewer motion although it always does when viewing true 3-D environment. This causes such unnatural effects as the perception of a rotating 3D object induced by sideways head motion, even if the object is static (e.g., single stereo pair). The motion-parallax conflict can be alleviated by view adaptation, i.e., the delivery to each eye of suitable views depending on current viewer position. This can be achieved by standard stereoscopic displays in combination with head tracking; in response to head position change suitable view is presented on the screen. The computation of views is not a problem for computer-rendered data, however it is difficult for camera-acquired images.³

An interesting alternative is a passive mechanism used in autostereoscopic displays where switching of views in time (temporal multiplexing) is replaced by spatial multiplexing, although at the expense of reduced spatial resolution. An early version of such automultiscopic (multiview autostereoscopic) display based on LCD technology

was proposed by Philips Research Laboratories.⁴ Recently, two commercial automultiscopic displays have been offered. Stereographics Corp. has developed a line of 9-view lenticular displays, called the *SynthaGram*TM,⁵ while 4-D Vision GmbH of Germany – an 8-view display based on wavelength-selective filter arrays.⁶ Both technologies have been demonstrated on LCD and plasma screens.

We have been experimenting with a 9-view 18-inch LCD version of the *SynthaGram*TM in which a lenticular sheet is used to create multiple viewing zones. A straightforward extension of the two-view lenticular autostereoscopic principle to M views, would result in an M -fold loss of the horizontal resolution. In order to circumvent the imbalance between the horizontal and vertical resolution, Stereographics Corp. applied a small tilt to the lenticules. In this way a better trade-off between the vertical and horizontal resolution loss in each view can be achieved.

The M -fold loss of the horizontal resolution in the case of purely vertical lenticules is the result of horizontal M -fold subsampling of each view in order to multiplex the M views on the screen. Clearly, this subsampling should be preceded by a suitable lowpass pre-filtering in order to avoid aliasing. In this case, a simple horizontal lowpass filtering with the bandwidth of $[-\pi/M, \pi/M]$ would suffice. The situation, however, is more complicated in the case of diagonal lenticules. The subsampling of each view is no more orthogonal (retained samples are not aligned vertically between rows), and more complex sampling models need to be considered.

Based on the multidimensional sampling theory,⁷ we study in this paper multiview sampling models based on lattices. These models can approximate pixel positions on lenticular automultiscopic displays thus leading to optimal anti-alias filters. We report results for a separable approximation to non-separable 2-D anti-alias filters based on the assumption that the lenticule slant is small. We carry out experiments on a variety of images, and different filter bandwidths. We observe that the theoretically-optimal bandwidth is too restrictive; aliasing artifacts disappear, but some image details are lost as well. Somewhat wider bandwidths result in images with almost no aliasing and largely preserved detail. For subjectively-optimized filters, the improvements, although localized, are clear and enhance the 3-D viewing experience.

2. REVIEW OF M-D SAMPLING THEORY

Automultiscopic displays convey 3-D information by means of multiplexing a number of views on the screen, and by employing a light-path selective mechanism such as lenticular sheet, parallax barrier, etc. At a specified distance from the screen, a viewer can see largely different sets of pixels by each eye due to the difference in angles formed by lines passing from each eye's optical center to a given screen point, thus inducing 3-D perception. If the display employs M views, and each original view is captured at full-screen resolution (e.g., 1280×1024), then in order to multiplex all views on the screen some form of $M:1$ subsampling must be applied to intensities/color of each view. As it will turn out in the next section, this subsampling is, in general, non-orthogonal (non-rectangular) and, therefore, cannot be described by two 1-D (one-dimensional) sampling processes, one horizontally and one vertically. A more general sampling structure is required in order to model the multiplexing process employed in automultiscopic displays such as the *SynthaGram*TM.

Below, we briefly review sampling theory of multidimensional (M-D) signals developed by Dubois.⁷ Let $u_c(\mathbf{x})$ be an N -dimensional continuous signal, i.e., $u_c \in R$ and $\mathbf{x} \in R^N$. For $N=2$ one can think of u_c as light distribution on a 2-D sensor at a single time instant, while for $N=3$ – the same distribution but as a function of time as well. Let the vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N\}$ form a basis, not necessarily orthogonal, of R^N . Then, *lattice* $\Lambda \subset R^N$ is defined as a set of discrete points in R^N formed by all linear combinations of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N$ with integer coefficients, i.e., $\Lambda = \{\mathbf{w} : \mathbf{w} = n_1\mathbf{v}_1 + \dots + n_N\mathbf{v}_N, n_i \in Z, i = 1, \dots, N\}$, where Z is the set of all integers. The set $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N\}$ is called a *basis* of lattice Λ . Clearly, $\mathbf{x} \in \Lambda$ can be represented as a linear combination of basis vectors: $\mathbf{x} = n_1\mathbf{v}_1 + \dots + n_N\mathbf{v}_N$. Alternatively, $\mathbf{x} = \mathbf{V}\mathbf{n}$ where $\mathbf{V} = [\mathbf{v}_1|\mathbf{v}_2|\dots|\mathbf{v}_N]$ is called the *sampling matrix* of Λ . Note that the columns of \mathbf{V} are formed by basis vectors of Λ .

We introduce now two concepts that are essential for frequency-domain representation of M-D signals. \mathcal{P} is a *unit cell* of lattice $\Lambda \in R^N$ if: 1) $\mathcal{P} \subset R^N$; 2) $\bigcup_{\mathbf{x} \in \Lambda} (\mathcal{P} + \mathbf{x}) = R^N$; and 3) $(\mathcal{P} + \mathbf{x}) \cap (\mathcal{P} + \mathbf{y}) = \emptyset$ for $\mathbf{x} \neq \mathbf{y}, \mathbf{x}, \mathbf{y} \in \Lambda$. Two important examples of unit cells are: parallelepiped and Voronoi cell. Then, the set of vectors $\mathbf{y} \in R^N$ such that $\mathbf{y}^T \mathbf{x} \in Z \forall \mathbf{x} \in \Lambda$ is called a *reciprocal lattice*, denoted Λ^* , with respect to the lattice Λ (Z is again the set of all integers).

Let $u(\mathbf{x})$ be an M-D signal defined over lattice Λ , i.e., with $\mathbf{x} \in \Lambda$. The Fourier transform of $u(\mathbf{x})$ is defined as follows:

$$U(\mathbf{f}) = \sum_{\mathbf{x} \in \Lambda} u(\mathbf{x}) e^{-j2\pi \mathbf{f} \cdot \mathbf{x}}, \quad \mathbf{f} \in R^N.$$

It can be easily verified that for $\mathbf{r} \in \Lambda^*$, $U(\mathbf{f} + \mathbf{r}) = U(\mathbf{f})$, i.e., the spectrum of signal sampled on lattice Λ is periodic with periodicity defined by the reciprocal lattice Λ^* . This is analogous to spectral periodicity at integer multiples of $2\pi f_s$ in the 1-D case, where f_s denotes 1-D sampling frequency.

We are ready now to see the frequency-domain consequences of M-D sampling. Let the samples of u be drawn from the continuous signal u_c as follows: $u(\mathbf{x}) = u_c(\mathbf{x})$, $\mathbf{x} \in \Lambda \subset R^N$. Then, it can be shown⁷ that the Fourier transform of the sampled signal u is:

$$U(\mathbf{f}) = \frac{1}{|\det V|} \sum_{\mathbf{r} \in \Lambda^*} U_c(\mathbf{f} + \mathbf{r}),$$

where U_c is the Fourier transform of the continuous signal u_c . The immediate consequence of the above summation is that a spectral overlap, known as *aliasing*, may occur after sampling if U_c is not sufficiently band-limited. If aliasing occurs, the original continuous signal u_c cannot be reconstructed from u without distortion. Spatial aliasing is often visible as jagged edges in still images ($N=2$), while temporal aliasing can be noticed in motion sequences as the opposite direction of sprocket wheel rotation ($N=3$). It has been shown,⁷ that in order to assure no aliasing in the sampled signal $u(\mathbf{x})$ it is sufficient (although not necessary) to limit the spectral support of the continuous signal to the Voronoi cell \mathcal{P}_{Λ^*} of the reciprocal lattice Λ^* , i.e., $U_c(\mathbf{f}) = 0$ for $\mathbf{f} \notin \mathcal{P}_{\Lambda^*}$.

In the case of automultiscopic displays, we assume that the M views have been properly acquired and contain no aliasing. These views are then subsampled by factor of $M:1$ in order to allow multiplexing on the screen. If we can model this subsampling process by lattice Λ then, in order that the subsampled views contain no aliasing, we must apply lowpass filtering with the passband limited to the Voronoi cell of Λ^* .

3. SAMPLING MODEL

A schematic representation of how 9 views are seen by a viewer in front of an automultiscopic lenticular display, such as the *SynthaGram*TM, is shown in Fig. 1. *RGB* components of pixels from each view can be assigned to screen *RGB* components (sub-pixels) depicted in Fig. 1 in a variety of ways. For example, a pixel component from one view can be assigned to one screen sub-pixel (component) only, or it can be divided between several neighboring sub-pixels.




Figure 1. Schematic representation of lenticule positions with respect to *RGB* sub-pixels on an LCD display, and fictitious representation of 9 views visible from different angles.

Fig. 2 shows an assignment of single-view intensities to pixels' *RGB* components similar to the one used in 9-view 18-inch *SynthaGram*TM monitor. The plot is at pixel resolution, with three components assigned to each

pixel (circles – R , squares – G , triangles – B). A circle, square or triangle present means that a value is assigned to the corresponding component at this pixel. Assignment for view #1 is shown in Fig. 2.a, and that for view #2 – in Fig. 2.b. Clearly, a complicated, non-orthogonal subsampling takes place that cannot be accurately modeled by independent horizontal and vertical sampling processes.

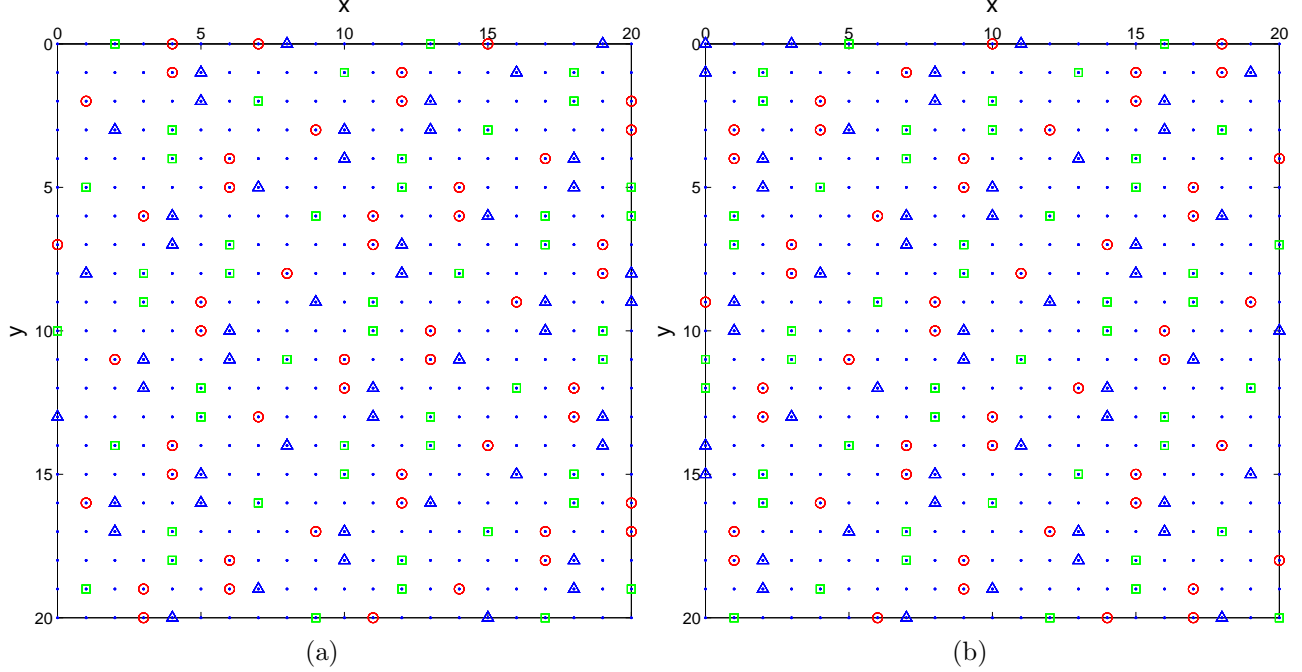


Figure 2. Simulation of an assignment of single-view intensities to RGB components of a 9-view 18-inch *SynthaGram*TM monitor from Stereographics Corp. (see acknowledgments): (a) view #1; and (b) view #2. The representation is at pixel resolution with circles denoting the red component (R), squares – green component (G), and triangles – blue component (B).

We propose to model the above view subsampling using the concept of lattice developed in Section 2. Let Γ be an orthonormal lattice (traditional rectangular sampling) on which pixels of one view are defined. As stated earlier, each pixel consists of three components RGB . In Fig. 2, lattice Γ is depicted by dots. We search for a lattice Λ with the sampling matrix \mathbf{V}_Λ , such that positions $\mathbf{x} \in \Lambda$ coincide as close as possible with the subsampled locations of one component of one view (e.g., circles, squares or triangles in Fig. 2). Clearly, we can only find an approximate lattice Λ since locations of one component in Fig. 2 do not constitute a lattice. A more accurate model could be achieved by a union of shifted lattices, but this is beyond the scope of this paper.

In order to find the sampling matrix \mathbf{V}_Λ , we carry out the following minimization:

$$\min_{\mathbf{V}_\Lambda} \sum_{\mathbf{y} \in \Gamma} |\xi_\Lambda(\mathbf{y}) - \xi_C(\mathbf{y})|, \quad (1)$$

where ξ_Λ is an indicator function for Λ , i.e.,

$$\xi_\Lambda(\mathbf{y}) = \begin{cases} 1 & \text{if } \mathbf{y} = \text{nint}(\mathbf{x}) \text{ for some } \mathbf{x} = [x_1 \ x_2]^T \in \Lambda, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

with $\text{nint}(\mathbf{x}) = [\text{nint}(x_1) \ \text{nint}(x_2)]^T$, and $\text{nint}(a)$ being the nearest integer of $a \in R$. Similarly, $\xi_C(\mathbf{y})$ is an indicator function for component C , i.e., either R (circle), G (rectangle) or B (triangle):

$$\xi_C(\mathbf{y}) = \begin{cases} 1 & \text{if } \mathbf{y} \text{ coincides with location of component } C, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The above minimization searches for parameters of the sampling matrix \mathbf{V}_Λ such that locations of the lattice Λ maximally coincide with those of one “interzigged” component. We implemented this minimization by means of hierarchical exhaustive search, and for the subsampling patterns of one component from Fig. 2 we obtained:

$$\mathbf{V}_\Lambda = \begin{bmatrix} 1.00 & 0.00 \\ 8.79 & 13.89 \end{bmatrix}. \quad (4)$$

Note that entries in the above matrix are expressed with respect to those of the unit matrix \mathbf{V}_Γ . To verify the accuracy of our sampling model, Fig.3 shows lattice Λ for the above sampling matrix overlaid onto R and G components of views #1 and #2. Note an accurate overlay for isolated component sub-pixels and reasonable in-between positions for sub-pixel pairs. For visualization reasons, Fig.3 shows only a tiny portion of the full-image sampling grid, but we have confirmed that similar alignment occurs in the remaining parts of the grid.

As discussed in Section 2, in order to determine specifications of the anti-alias pre-filter for lattice Λ we need to consider its reciprocal lattice Λ^* . In Fig. 4 both the lattice Λ and its reciprocal Λ^* are shown as bullets, whereas the orthonormal lattice Γ and its reciprocal Γ^* – as triangles. Recall that reciprocal lattice locations in the frequency domain (Fig. 4) denote where the spectral replications occur. Clearly, since Λ^* is denser than Γ^* , the maximum support of a signal spectrum that would not induce aliasing can be larger in the case of Γ than in the case of Λ . This is better illustrated in Fig. 5.a where the reciprocal lattices Λ^* and Γ^* are shown together with their Voronoi cells (shaded). Again, to avoid aliasing when subsampling from Γ to Λ , the signal defined on Γ must be pre-filtered by a lowpass filter with the passband limited to the Voronoi cell of the reciprocal lattice Λ^* (the diamond-shaped shaded area in the center).

The shape of the Voronoi cell for Λ^* identifies the passband boundary of an ideal 2-D anti-alias pre-filter. However, such a filter is not separable into two 1-D filters, and thus its design and implementation are far from trivial. In the next section, we consider an approximation to these filter specifications that, although not very accurate, will give us an indication as to what visual improvements can be expected from pre-filtering.

4. ORTHOGONAL APPROXIMATION OF THE SAMPLING MODEL

Specifications of the anti-alias pre-filter discussed above are difficult to attain in a practical filter. First, the passband shape is not separable. Secondly, the transition-band width is zero to assure no aliasing and no loss of information. One must approximate the ideal pre-filter, and it is a design choice how close the approximation should be to the ideal specifications. Since for now we are interested only in validating the concept of pre-filtering when multiplexing views, we opt for a separable rectangular passband (aligned with frequency axes f_x, f_y), as shown in Fig. 5.b. The remaining issue is to decide as to the width of the rectangular passband. In order to assure no aliasing in the subsampled data, the passband must be small enough not to overlap when replicated on Λ^* . It turns out that a square passband with 1/4-bandwidth horizontally and vertically is a simple and sufficiently accurate choice. Although this filter permits aliasing at frequencies located in the square’s corners, we consider these effects negligible for now. Note that the 1/4-bandwidth means 1/4 of the Nyquist rate. Since all the lattice plots in this paper are normalized to the orthonormal lattices Γ and Γ^* , the passband of a 1/4-band filter is limited by ± 0.125 in both directions.

It is interesting to observe that if the above 1/4-band square-shaped passband is considered to characterize an optimal anti-alias pre-filter for some lattice, this lattice must be orthogonal (square shape of the Voronoi cell of the reciprocal lattice). Fig. 6 shows this orthogonal lattice Λ_o and its reciprocal lattice Λ_o^* . This is why we call Λ_o an orthogonal approximation to the sampling model.

5. EXPERIMENTAL RESULTS

5.1. Design of optimal anti-alias pre-filter

Since the 1/4-band square-shaped passband is separable, we need to design only a 1-D prototype filter. We used the Parks-McClellan algorithm for optimum FIR filter design implemented in *Matlab* through the `remez` function. Table 1 shows specifications of three filters with different passbands (1/3-band, 1/4-band, and 1/5-band) but the same transition band, in order to assure the same passband magnitude error. We will evaluate

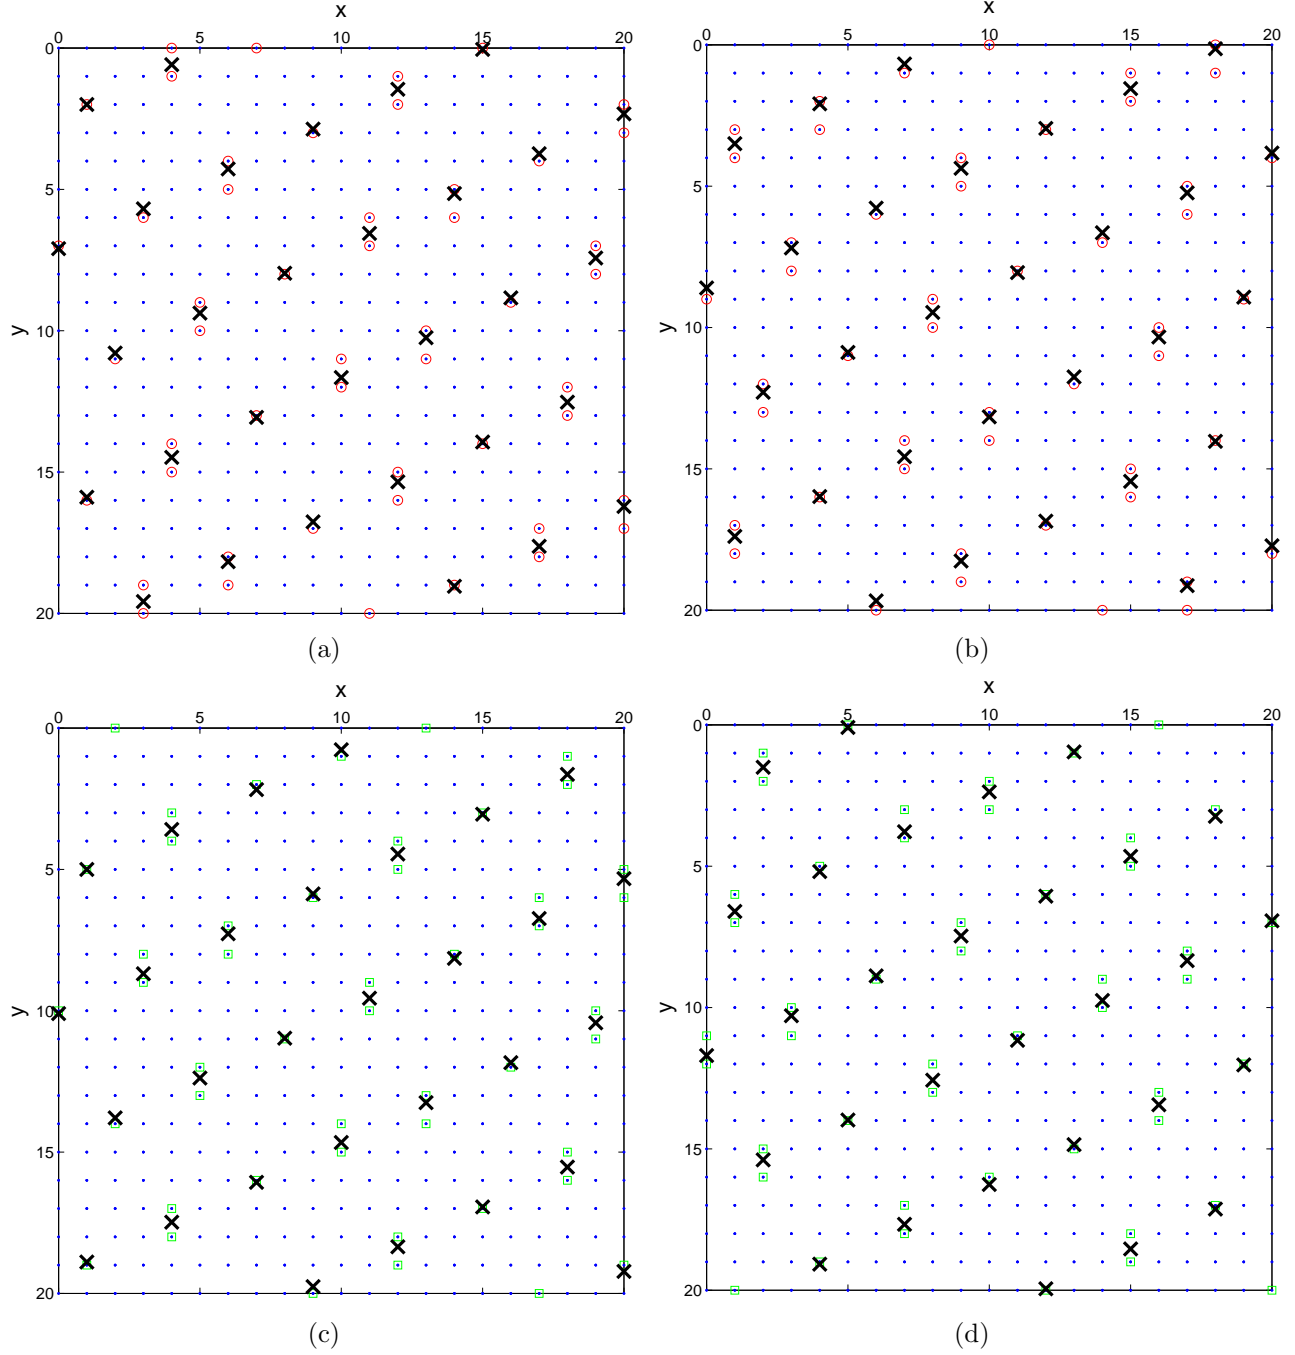


Figure 3. Lattice Λ (\times) obtained by the minimization (1) and overlaid (with a suitable shift) onto the assignments of red and green components from Fig. 2: (a) red component – view #1; (b) red component – view #2; (c) green component – view #1; and (d) green component – view #2.

later the impact of filter bandwidths on the amount of aliasing present and loss of resolution. As can be seen, filter specifications have been carefully selected so that the filter's magnitude response drops by 3dB at exactly 1/5-th, 1/4-th and 1/3-rd of the Nyquist rate, respectively.

Fig. 7 shows the 1-D magnitude and impulse responses of the designed filters, and also the 2-D magnitude responses of the effective 2-D filter implemented through the convolution of horizontal and vertical 1-D filters. Note the narrow transition band of the magnitude responses, and square shape of the passbands.

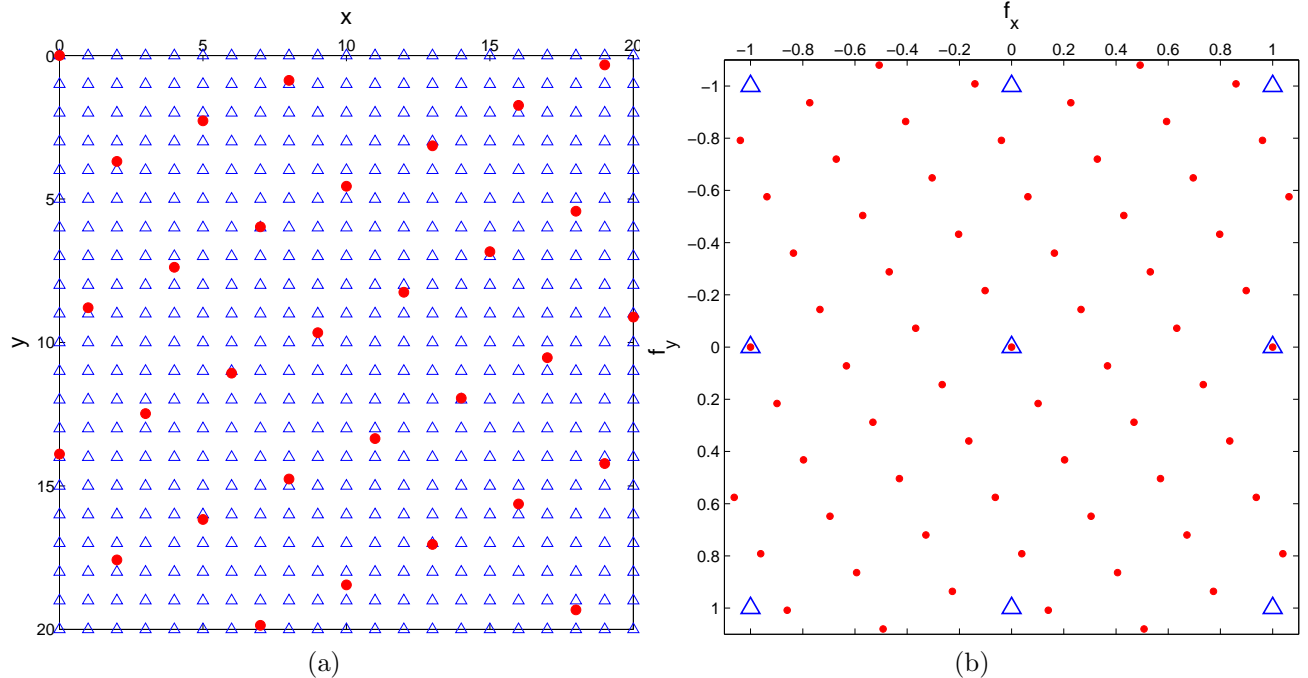


Figure 4. Subsampling model for one component of a single view: (a) lattice Λ overlaid onto Γ , and (b) reciprocal lattice Λ^* overlaid onto Γ^* ($\bullet = \Lambda, \Lambda^*$, and $\triangle = \Gamma, \Gamma^*$).

Table 1. Specifications and effective bandwidths of three 1-D FIR anti-alias pre-filters (normalized Nyquist rate = 1.0).

	# coeff.	desired passband	specified passband/stopband	effective passband (-3dB)
1/5-band filter	51	0 – 0.2000	0 – 0.167 / 0.257 – 1	0 – 0.1995
1/4-band filter	51	0 – 0.2500	0 – 0.217 / 0.307 – 1	0 – 0.2495
1/3-band filter	51	0 – 0.3333	0 – 0.301 / 0.391 – 1	0 – 0.3333

5.2. Importance of bandwidth limitation before subsampling

We first demonstrate the importance of image pre-filtering prior to subsampling. Figs. 8 and 9 show images *Wedding* and *Julie* (see acknowledgments) subsampled by 4 both horizontally and vertically without and with 1/4-band pre-filtering. Note graininess of the pavement and jagged edges (shadow of the lamp post, columns in the ruins) in *Wedding* with no pre-filtering, and again graininess of the wall and jagged edges on the hat in *Julie*. These effects are very clear when viewed on a computer monitor, but are less obvious in print due to the reduced dynamic range reproduced. Differences between images processed by the 1/4-band filter and the other two filters are subtle, and do not reproduce well in print. Clearly, lack of pre-filtering introduces distortion.

5.3. Subjective viewing of multiview images on the *SynthaGram*TM monitor

We have applied the designed filters to all 9 views of *Wedding* and *Julie*, and subsequently multiplexed them using the “interzigging” software from Stereographics Corp. Viewed on an 18-inch LCD 1280×1024 *SynthaGram*TM monitor the pre-filtered images offered clear although subtle improvements compared to the non-filtered ones. In *Wedding*, we observed the following improvements: eliminated false colors (mostly green and red dots) on sharp edges, improved detail on the bouquet of flowers, no “dirty-window” distortion on the metal railing under arches of the ruins, and no false colors on the heap of rocks under the arches to the left of the photographer. We have observed fewer improvements in *Julie*: eliminated false colors on the stripes of the hat, and improved edge detail, e.g., on the rim of the hat. The smaller improvement for *Julie* is due to the fact that the image contains

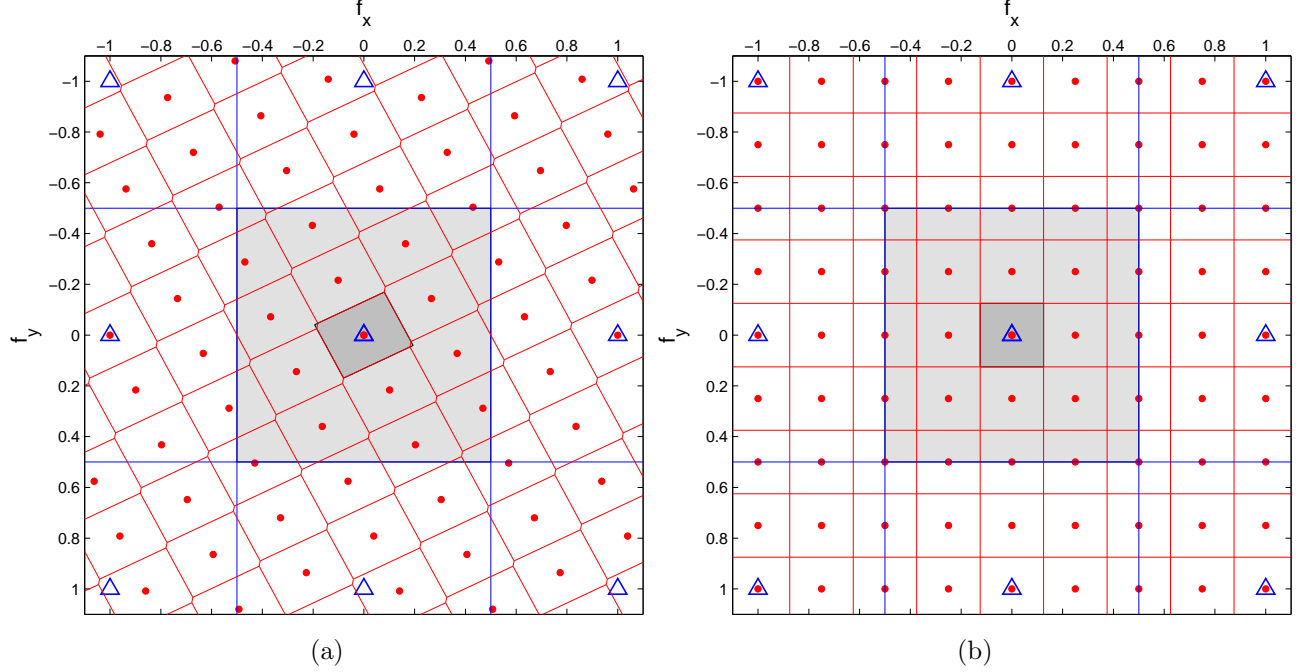


Figure 5. Reciprocal lattices: (a) Λ^* , and (b) Λ_o^* overlaid onto Γ^* together with their Voronoi cells (shaded areas at the center) and their repetitions. ($\bullet = \Lambda^*$ in (a) and Λ_o^* in (b); $\triangle = \Gamma^*$.)

less high-frequency content, in particular fewer high-contrast edges. We have observed similar improvements on other images, both camera-acquired and computer-rendered.

The observed improvements are exclusively in high-frequency image areas, thus confirming the importance of bandwidth limitation. Certainly a delicate balance between aliasing suppression and resolution loss needs to be found. We observed that 1/5-band filtering offered no more improvements than the 1/4-band filtering, while significantly reducing detail in high-frequency areas in both images. On the other hand, 1/3-band filtering, that theoretically does not offer sufficient bandwidth limitation, resulted only in a marginally increased aliasing but markedly improved detail compared to the 1/4-band filter. This suggests that further improvement might be possible if a more accurate approximation of filter’s passband to the Voronoi cell \mathcal{P}_Λ (Fig. 5.a) can be found.

Another interesting observation we made is that aliasing distortions are particularly visible at high-contrast edges, such as the lower arches in *Wedding*. We believe that the distortion is not solely due to a single-view aliasing component but also to the multiplexing algorithm. In particular, if a screen pixel on the dark side of a sharp high-contrast boundary contains *RGB* components from different views, then it is likely that a component from another view is located on the bright edge side, thus causing edge “coloring”. A remedy for this might be an adaptive multiplexing scheme that accounts for local contrast.

6. CONCLUSIONS

We have presented an analysis of 2-D sampling of views in a lenticular automultiscopic 3-D display. We have found parameters of a lattice that closely approximates this sampling pattern. Since this sampling requires non-separable 2-D filtering, we have approximated the band-limitation by a separable square-shaped 2-D filter that is equivalent to assuming an orthogonal sampling model. We have designed 1-D 51-coefficient zero-phase FIR filters meeting the orthogonal subsampling specifications. Applied to all views prior to multiplexing, these filters resulted in subtle although clear improvements on edges and in high-detail areas.

As we have pointed out, a slightly wider filter resulted in improved detail without significantly increased aliasing. This suggests that a better optimization of filter bandwidth might be beneficial. To this effect, we are planning to study non-orthogonal sampling that leads to non-separable 2-D anti-alias filters. One possible path

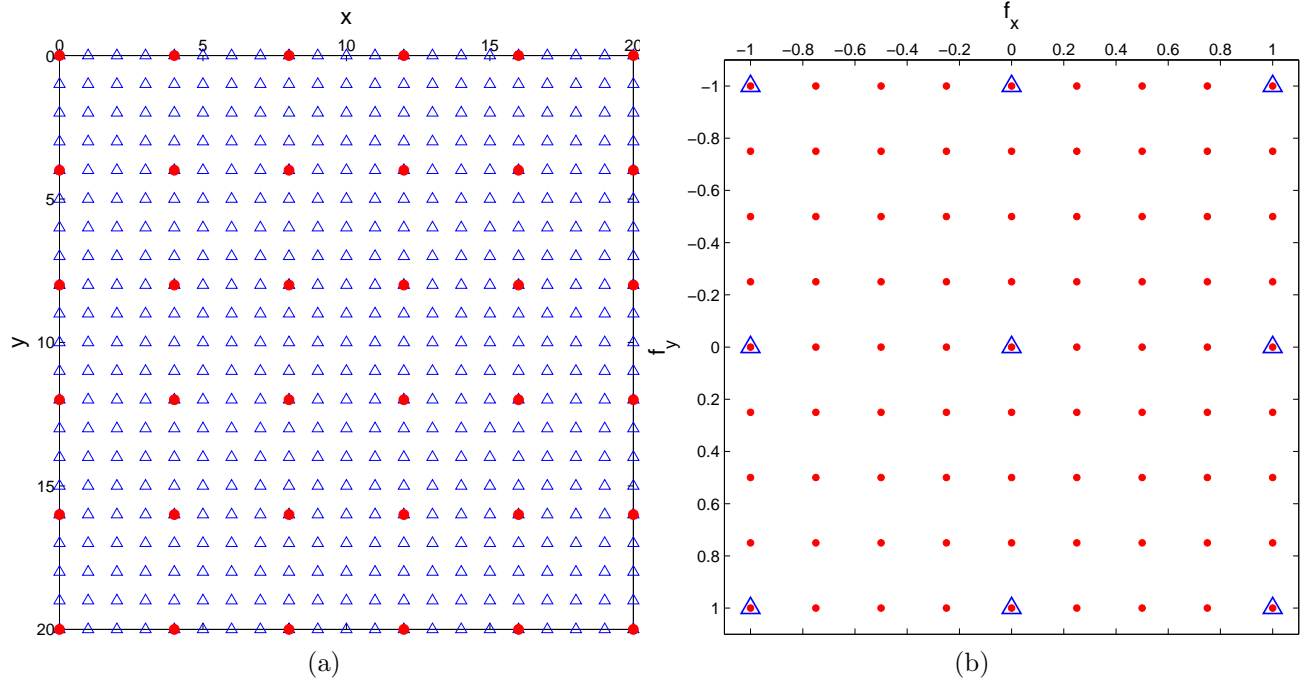


Figure 6. Orthogonal approximation of the subsampling model from Fig. 4 for one component of a single view: (a) lattice Λ_o overlaid onto Γ , and (b) reciprocal lattice Λ_o^* overlaid onto Γ^* ($\bullet = \Lambda_o, \Lambda_o^*$, and $\triangle = \Gamma, \Gamma^*$).

to a straightforward design and implementation of non-orthogonal filtering is the Fourier transform. Efficient implementation of the Fourier transform through the FFT (Fast Fourier Transform), and the relative ease of specification of arbitrary passband shape in the frequency domain, provide a clear alternative to complex design and implementation issues in the spatial domain.⁸

Finally, a careful inspection of the lattice Λ overlaid on individual components in Fig. 3 reveals that a more accurate sampling model would be possible with a union of shifted lattices, so that *all* component positions (circles, squares or triangles) would coincide with samples of such a model. This would require the so-called *union of cosets* sampling⁷ and call for even more sophisticated 2-D filtering.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (NSF) under Grant ECS-0219224. The authors would like to thank Stereographics Corporation, and especially Mr. Mark Feldman, for providing technical specifications of the *SynthaGram*TM monitor (18") and of the "interzigging" process that were used in the course of this research. They would also like to thank Mr. Lenny Lipton for his permission to use images *Wedding* and *Julie* to carry out experiments.

REFERENCES

1. E. Dubois, "A projection method to generate anaglyph stereo images," in *Proc. IEEE Int. Conf. Acoustics Speech Signal Processing*, **3**, pp. 1661–1664, May 2001.
2. J. Konrad, B. Lacotte, and E. Dubois, "Cancellation of image crosstalk in time-sequential displays of stereoscopic video," *IEEE Trans. Image Process.* **9**, pp. 897–908, May 2000.
3. J. Konrad, "View reconstruction for 3-D video entertainment: issues, algorithms and applications," in *Proc. Int. Conf. on Image Process. and its Applications*, pp. 8–12, July 1999.
4. C. Van Berkel, A. Franklin, and J. Mansell, "Design and applications of multiview 3D-LCD," in *Proc. SID Euro-Display '96*, pp. 109–112, 1996.

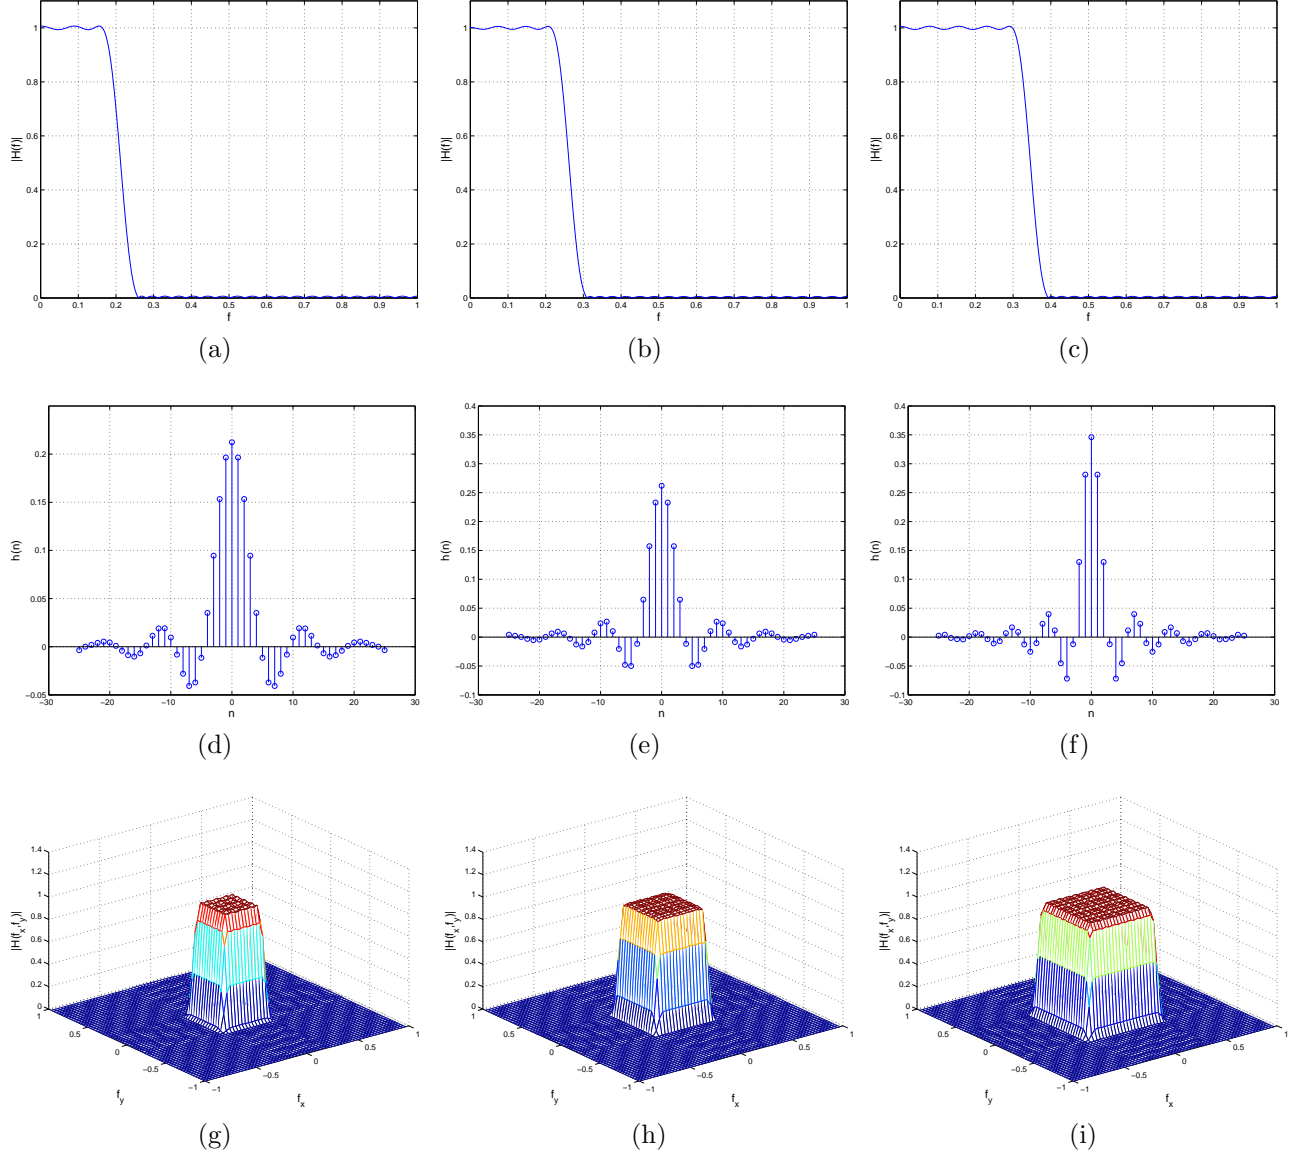


Figure 7. Magnitude and impulse responses of 1/5-, 1/4-, and 1/3-band anti-alias filters used in experiments. (a-c) 1-D magnitude response; (d-f) 1-D impulse response; and (g-i) 2-D magnitude response of separable 2-D filter. Normalized frequency of 1.0 corresponds to the Nyquist rate.

5. L. Lipton and M. Feldman, “A new stereoscopic display technology: The SynthaGram,” in *Proc. SPIE Stereoscopic Displays and Applications*, **4660**, pp. 229–235, Jan. 2002.
6. A. Schmidt and A. Grasnack, “Multi-viewpoint autostereoscopic displays from 4D-Vision,” in *Proc. SPIE Stereoscopic Displays and Applications*, **4660**, pp. 212–221, Jan. 2002.
7. E. Dubois, “The sampling and reconstruction of time-varying imagery with application in video systems,” *Proc. IEEE* **73**, pp. 502–522, Apr. 1985.
8. J. Radecki, J. Konrad, and E. Dubois, “Design of multidimensional finite-wordlength FIR and IIR filters by simulated annealing,” *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.* **42**, pp. 424–431, June 1995.



(a)



(b)

Figure 8. Illustration of 2-D subsampling artifacts for image *Wedding*: (a) 4:1 subsampling with no pre-filtering; (b) 4:1 subsampling with 1/4-band pre-filtering (Fig. 7.b).



(a)



(b)

Figure 9. Illustration of 2-D subsampling artifacts for image *Julie*: (a) 4:1 subsampling with no pre-filtering; (b) 4:1 subsampling with 1/4-band pre-filtering (Fig. 7.b).