

Multidimensional Directional Filter Banks and Surfacelets

Yue M. Lu, *Student Member, IEEE*, and Minh N. Do, *Member, IEEE*

Abstract—In 1992, Bamberger and Smith proposed the directional filter bank (DFB) for an efficient directional decomposition of 2-D signals. Due to the nonseparable nature of the system, extending the DFB to higher dimensions while still retaining its attractive features is a challenging and previously unsolved problem. We propose a new family of filter banks, named NDFB, that can achieve the directional decomposition of arbitrary N -dimensional ($N \geq 2$) signals with a simple and efficient tree-structured construction. In 3-D, the ideal passbands of the proposed NDFB are rectangular-based pyramids radiating out from the origin at different orientations and tiling the entire frequency space. The proposed NDFB achieves perfect reconstruction via an iterated filter bank with a redundancy factor of N in N -D. The angular resolution of the proposed NDFB can be iteratively refined by invoking more levels of decomposition through a simple expansion rule. By combining the NDFB with a new multiscale pyramid, we propose the *surfacelet* transform, which can be used to efficiently capture and represent surface-like singularities in multidimensional data.

Index Terms—Directional decomposition, directional filter banks (DFBs), filter design, high-dimensional transforms, surfacelets.

I. INTRODUCTION

WITH the growing capabilities of modern computers and imaging devices, high-resolution 3-D and even higher dimensional volumetric data are increasingly available in a wide gamut of applications, including biomedical imaging, seismic imaging, extragalactic astronomy, computer vision, and video processing and compression. To efficiently analyze and represent such huge amount of data, we need to create and employ new tools from various fields of engineering, including signal processing. In this paper, we present a new set of tools, namely the N -dimensional directional filter banks (NDFBs) and surfacelets, that can capture and represent signal singularities lying on smooth surfaces. Such singularities are often observed in 3-D medical images, where the images are mostly smooth except on some boundary surfaces, and in video signals, in which moving objects carve out smooth surfaces in the 3-D spatial/temporal space.

Manuscript received December 19, 2005; revised October 29, 2006. This work was supported by the U.S. National Science Foundation under Grant CCR-0237633 (CAREER). The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Ljubisa Stankovic.

Y. M. Lu is with the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: yuelu@uiuc.edu).

M. N. Do is with the Department of Electrical and Computer Engineering, the Coordinated Science Laboratory, and the Beckman Institute, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: minhdo@uiuc.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2007.891785

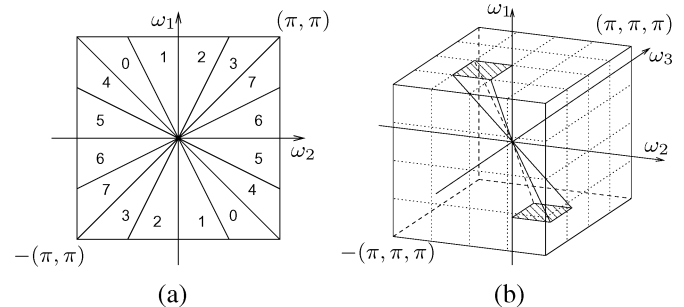


Fig. 1. (a) Frequency partitioning of the directional filter bank with three levels of decomposition. There are $2^3 = 8$ real wedge-shaped frequency bands. (b) Frequency partitioning of the proposed NDFB in 3-D. The ideal passbands of the component filters are rectangular-based pyramids radiating out from the origin at 3×2^l ($l \geq 0$) different orientations and tiling the entire frequency space.

For 2-D signals, the similar problem of capturing singularities along smooth curves has been extensively studied. Without claiming to be exhaustive, we would like to mention a few examples, including the steerable pyramid [1], the directional filter bank [2], 2-D directional wavelets [3], curvelets [4], complex wavelets [5], [6], contourlets [7], bandelets [8], and shearlets [9].

Among all these 2-D representations, one approach of particular interest to us is the directional filter bank (DFB), which was originally proposed by Bamberger and Smith [2] and subsequently improved by several authors [10]–[16]. The DFB is efficiently implemented via an l -level tree-structured decomposition that leads to 2^l subbands with wedge-shaped frequency partitioning as shown in Fig. 1(a). Meanwhile, the DFB is a nonredundant transform, and offers perfect reconstruction, i.e., the original signal can be exactly reconstructed from its decimated channels. The directional-selectivity and efficient structure of the DFB makes it an attractive candidate for many image processing applications. By combining the DFB with the Laplacian pyramid, Do and Vetterli [7] constructed the *contourlets*, which provides a directional multiresolution transform for sparse image representation.

The major contribution of this work is extending the DFB to higher dimensions. For example, in 3-D, we want to achieve the frequency partitioning as shown in Fig. 1(b), where the ideal passbands of the component filters are rectangular-based pyramids radiating out from the origin at different orientations and tiling the entire frequency space. We can see this is a natural extension from the wedge-shaped frequency partitioning in 2-D.

Filters with this type of pyramid-shaped frequency supports can be used in selective filtering of plane wave signals based on their directions of arrival [17], as well as efficiently capturing

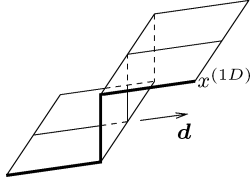


Fig. 2. Plane wave with a step profile $x^{(1D)}$ along direction \mathbf{d} .

signal singularities (discontinuities) living on surfaces with different normal directions. For example, consider an N -D continuous plane-wave signal $x(\mathbf{t})$, completely determined by some 1-D function $x^{(1D)}(\cdot)$, i.e.,

$$x(\mathbf{t}) = x^{(1D)}(\mathbf{d}^T \cdot \mathbf{t}),$$

where $\mathbf{t} = (t_1, \dots, t_N)^T$, and \mathbf{d} is a unit norm directional vector. As illustrated in Fig. 2, when $x^{(1D)}$ is a step function, the plane wave $x(\mathbf{t})$ then represents an ideal singularity across a surface of co-dimension 1 in \mathbb{R}^N . We can verify that the region of support of the Fourier transform of $x(\mathbf{t})$ lies on a straight line radiating out from the origin and oriented in direction \mathbf{d} . If we decompose a sampled version of $x(\mathbf{t})$ using NDFB, then most of the signal energy will be captured by only a few pyramid-shaped subbands whose directions are aligned with \mathbf{d} .

Unlike the separable wavelets [18]–[20], whose multidimensional generalizations are simply the tensor products of their 1-D counterparts, the DFB has a much more involved nonseparable construction. Extending the DFB to higher dimensions while still retaining its various attractive features is a challenging and, to our best knowledge, previously unsolved problem. In this paper, we propose a new family of filter banks, named NDFB, with the following distinctive properties.

- 1) **Directional decomposition.** The NDFB decomposes N -D ($N \geq 2$) signals into directional subbands, as shown in Figs. 1(a) and (b) for 2-D and 3-D, respectively. In general, the ideal passbands of the NDFB in higher dimensions are hypercube-based hyperpyramids radiating out from the origin.
- 2) **Construction.** The NDFB has an efficient tree-structured implementation using iterated filter banks.
- 3) **Angular resolution.** The number of directional subbands can be increased by iteratively invoking more levels of decomposition through a simple expansion rule. In general, there can be $N \times 2^l$ ($l \geq 0$) different directional subbands in the N -D case.
- 4) **Perfect reconstruction.** The original signal can be exactly reconstructed from its transform coefficients in the absence of noise or other processing.
- 5) **Small redundancy.** The NDFB is N -times expansive in the N -dimensional case.

The quest for extending the DFB to higher dimensions has been considered by several researchers in the past. In [10], Bamberger proposed a 3-D subband decomposition scheme implemented by applying the checkerboard filter banks separately along two orthogonal signal planes followed by a 2-D DFB decomposition on one of the planes. However, the resulting passband shapes are 3-D triangular prisms and do not correspond to

a single dominant direction. Meanwhile, the angular resolution can only be refined along one of the axes. In [12], Park proposed a 3-D velocity selective filter bank by applying two 2-D DFBs separately along two signal planes. The resulting frequency partitioning is similar to that of NDFB. However, that construction has a redundancy factor of 2^l for l -levels of decomposition. We would like to emphasize that our proposed NDFB has a relatively small redundancy ratio of 3 in 3-D, independent of the number of decomposition levels. More importantly, our NDFB construction can be easily generalized to arbitrary high dimensions.

The outline of the paper is as follows. We introduce the necessary notations in Section II, as well as review some concepts in multidimensional multirate systems. We then present our main results in two stages. Section III deals with the construction and properties of the proposed NDFB in 3-D. The generalization to arbitrary N -dimensional ($N \geq 2$) cases is given in Section IV. We discuss filter design and implementation issues in Section V. To efficiently capture local surface singularities with different sizes, in Section VI, we combine the NDFB with a multiscale decomposition and construct the *surfacelet* transform, whose basis images are local surface patches spanning a wide range of normal directions, spatial locations, and scales. We present experimental results in Section VII and conclude the paper in Section VIII with some discussions.

II. PRELIMINARIES

Since the major part of this work is about constructing a multidimensional filter bank, we start by reviewing some of the concepts and results in multidimensional multirate systems [19], [21], [22].

Notations: Throughout the paper, N represents the dimension of the signals. We are interested in cases when $N \geq 2$. We use lower-case letters, e.g., $x[\mathbf{n}]$ to denote N -D discrete signals, where $\mathbf{n} \stackrel{\text{def}}{=} (n_1, n_2, \dots, n_N)^T$ is an integer vector. The discrete-time Fourier transform of a multidimensional signal is defined as

$$X(\omega) = \sum_{\mathbf{n} \in \mathbb{Z}^N} x[\mathbf{n}] e^{-j\omega^T \mathbf{n}}.$$

In multidimensional multirate signal processing, the sampling operations are defined on lattices. A lattice in N -dimension is represented by an $N \times N$ nonsingular integer matrix \mathbf{M} as

$$\text{LAT}(\mathbf{M}) = \{\mathbf{m} : \mathbf{m} = \mathbf{M}\mathbf{n}, \mathbf{n} \in \mathbb{Z}^N\}.$$

For an \mathbf{M} -fold downsampling (or upsampling) operation, the input $x[\mathbf{n}]$ and the downsampling output $y_d[\mathbf{n}]$ (or the upsampling output $y_u[\mathbf{n}]$) are related by

$$y_d[\mathbf{n}] = x[\mathbf{M}\mathbf{n}] \text{ and } y_u[\mathbf{n}] = \begin{cases} x[\mathbf{M}^{-1}\mathbf{n}], & \text{if } \mathbf{n} \in \text{LAT}(\mathbf{M}) \\ 0, & \text{otherwise} \end{cases}$$

respectively. For the upsampling case, the Fourier transforms of the input and output are related by

$$Y_u(\omega) = X(\mathbf{M}^T \omega).$$

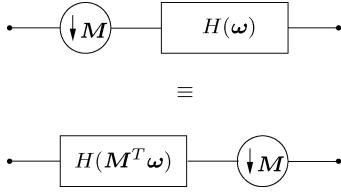


Fig. 3. Multidimensional multirate identity for interchanging the order of downsampling and filtering.

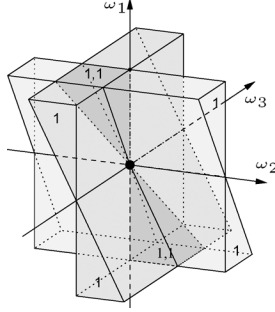


Fig. 4. (Dark region) Pyramid-shaped subband support is the intersection of (gray regions) two wedge-shaped supports.

Multirate identities [22] are often useful in analyzing multidimensional multirate systems. The identity for the analysis part of the filter bank is shown in Fig. 3; the one for the synthesis part can be inferred similarly. Downsampling by \mathbf{M} followed by filtering with a filter $H(\omega)$ is equivalent to filtering with the filter $H(\mathbf{M}^T \omega)$, which is obtained by upsampling $H(\omega)$ by \mathbf{M} , before downsampling.

III. NDFB IN 3-D

To simplify the exposition, we first discuss the proposed NDFB in 3-D. Once the construction in the 3-D case is clear, the generalization to arbitrary N -dimensional ($N \geq 2$) cases can easily be developed. We focus on the analysis part of the NDFB, since the synthesis part is exactly symmetric. Meanwhile, since we are mainly interested in the passband and stopband regions of the filters, we assume all the filters used in this section and Section IV are ideal, i.e., the filter frequency responses are indicator functions of the ideal passband supports. In Section V, we design real filters that approximate these conditions.

A. Key Ideas

Fig. 4 illustrates one of the key ideas of NDFB. The dark region in the figure shows the pyramid-shaped support in frequency that we want to get. We can see that the pyramid support can be obtained as an intersection of two wedge supports shown in gray regions. This observation leads to the idea of obtaining pyramid-shaped supports in 3-D DFB by a concatenation of two 2-D DFBs on appropriate dimensions.

Fig. 5(a) shows a wedge-shaped decomposition of the 3-D frequency spectrum. This decomposition can be achieved by applying a 2-D filter bank, e.g., the 2-D DFB with frequency decomposition as shown in Fig. 1(a), on every (n_1, n_2) -slice in the 3-D signal. For notations, we use

$$W_i^{(l_2)}(\omega_1, \omega_2), \quad 0 \leq i < 2^{l_2}$$

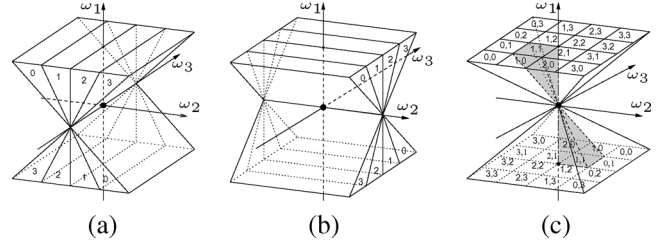


Fig. 5. (a) Ideal wedge-shaped frequency support of a 2-D filter operating along the (n_1, n_2) -plane. (b) The wedge-shaped support of a 2-D filter along the (n_1, n_3) -plane. (c) The ideal pyramid-shaped frequency decomposition.

to denote the ideal filter whose frequency support is on the i th wedge in Fig. 5(a). The superscript (l_2) indicates that there are 2^{l_2} wedge subbands (in this case, $l_2 = 2$) oriented at angles from -45° to 45° . The frequency variables ω_1 and ω_2 specify that the 2-D filters operate on (n_1, n_2) -slices. Similarly, we show in Fig. 5(b) the wedge-shaped frequency decomposition along the (ω_1, ω_3) plane. With the same notation above, we can use $W_j^{(l_3)}(\omega_1, \omega_3)$ ($0 \leq j < 2^{l_3}$) to represent the corresponding ideal subband filters.

Now, by taking the pairwise intersection of the wedge supports from Fig. 5(a) and (b), we can get 16 “thinner” pyramid supports, as shown in Fig. 5(c). In general, the hourglass-shaped region can be divided into $2^{l_2} \times 2^{l_3}$ ($l_2, l_3 \geq 0$) different square-based pyramids radiating out from the origin. As illustrated in Fig. 5(c), each subband can be indexed by a pair of integers (i, j) specifying the square base of the pyramid. We use

$$P_{i,j}^{(l_2, l_3)}(\omega_1, \omega_2, \omega_3), \quad 0 \leq i < 2^{l_2} \text{ and } 0 \leq j < 2^{l_3}$$

to represent the ideal pyramid filters.

With these notations, it is easy to check that (see Lemma 1 in Section IV)

$$P_{i,j}^{(l_2, l_3)}(\omega_1, \omega_2, \omega_3) = W_i^{(l_2)}(\omega_1, \omega_2) \cdot W_j^{(l_3)}(\omega_1, \omega_3) \quad (1)$$

for all $l_2, l_3 \geq 0$ and $0 \leq i < 2^{l_2}$, $0 \leq j < 2^{l_3}$. Usually, we say a multidimensional filter $F(\omega)$ is *separable*, when it can be written as the product of several 1-D filters, i.e., $F(\omega) = \prod_{i=1}^N F_i(\omega_i)$. Here, we extend this notion to *Kth-order generalized separability*, to describe those N -dimensional filters that can be represented as the product of several filters of dimensions up to K , with $0 < K < N$. We note that the pyramid filters $P_{i,j}^{(l_2, l_3)}(\omega_1, \omega_2, \omega_3)$ defined above are second-order generalized separable.

At this point, a natural question is: since the pyramid filters are the products of two wedge filters, can we achieve the pyramid-shaped frequency decomposition by directly applying two 2-D DFBs separately along the (n_1, n_2) and (n_1, n_3) signal planes? This idea was explored by Park in [12], but at a high price.

Here is why. The flow of operations by applying two 2-D DFBs sequentially along two signal planes can be illustrated as follows:

$$\begin{aligned} & \rightarrow \boxed{W_i^{(l_2)}(\omega_1, \omega_2)} \rightarrow \underbrace{(\downarrow m_1)}_{\text{along } n_1} \rightarrow \underbrace{(\downarrow m_2)}_{\text{along } n_2} \rightarrow \bullet \\ & \bullet \rightarrow \boxed{W_j^{(l_3)}(\omega_1, \omega_3)} \rightarrow \underbrace{(\downarrow m_1)}_{\text{along } n_1} \rightarrow \underbrace{(\downarrow m_2)}_{\text{along } n_3} \end{aligned} \quad (2)$$

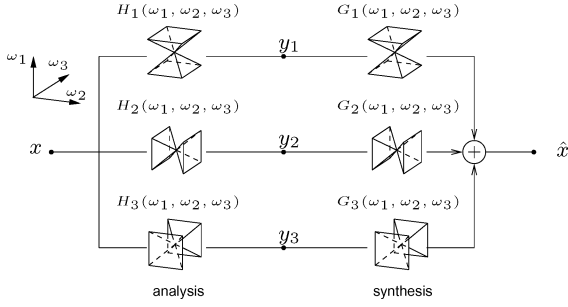


Fig. 6. First level of decomposition: three-channel undecimated filter bank in 3-D. The ideal frequency-domain supports of the component filters are hourglass-shaped regions, with their corresponding dominant directions aligned with the ω_1 , ω_2 , and ω_3 axes, respectively.

where $W_i^{(l_2)}(\omega_1, \omega_2)$ and $W_j^{(l_3)}(\omega_1, \omega_3)$ are two wedge-shaped filters from the two DFBs, respectively. In a critically sampled filter bank decomposition, filtering is always followed by downsampling. It can be shown [13] that the equivalent downsampling matrix for an l_2 -level DFB operating along the (n_1, n_2) plane is a diagonal matrix, implemented separately as downsampling in the n_1 dimension by $m_1 = 2$ followed by downsampling in the n_2 dimension by $m_2 = 2^{(l_2-1)}$. From the multirate identities introduced in Section II (Fig. 3), the downsampling by 2 along the n_1 dimension scrambles the wedge-shaped frequency decomposition provided by $W_j^{(l_3)}(\omega_1, \omega_3)$. Thus, it can be easily checked that the subsequent application of $W_j^{(l_3)}(\omega_1, \omega_3)$ will not provide the desired pyramid-shaped frequency decomposition as shown in Fig. 5(c).

To overcome this problem, Park [12] proposed to upsample and interpolate the decimated outputs of the first DFB to the original size before feeding them to the second DFB. With this step, the DFB essentially becomes a nonsampled filter bank, and, hence, the generalized separability described in (1) can be applied. However, this scheme leads to a highly redundant system (2^l times redundant for l levels of decomposition).

In the following, we propose a new filter bank structure, that can make use of the generalized separability property with a simple expansion rule, but without the excessive redundancy. Similar to (2), this new construction also contains the concatenation of two 2-D filter banks along two signal planes. However, a distinctive feature is that all the downsampling operation for the first 2-D filter bank is done on the n_2 dimension only (i.e., $m_1 = 1$). This ensures that the subsequent 2-D filter bank, which operates on the (n_1, n_3) dimensions, will not be affected by the downsampling of the previous filter bank.

B. First Level: The Hourglass Filter Bank

To obtain the first level of decomposition, we employ a three-channel undecimated filter bank shown in Fig. 6. This filter bank decomposes the 3-D frequency spectrum of the input signal into three hourglass-shaped subbands, with their dominant directions aligned with the ω_1 , ω_2 , and ω_3 axes, respectively.

Despite the redundancy it brings in, the undecimated hourglass filter bank in this step offers several important advantages over a decimated filter bank. As will be seen shortly, the undecimated 3-D filter bank allows subsequent levels of the NDFB to be implemented by using only 2-D filter banks. Meanwhile, it

is easier to design an undecimated filter bank with perfect reconstruction than a decimated one, since the former imposes a smaller set of constraints.

We should make one simplification before describing further levels of decomposition. By geometric symmetry, we only need to consider subsequent decomposition steps after the top branch of the hourglass filter bank, i.e., the y_1 subband in Fig. 6 whose dominant direction is along the ω_1 axis. The decomposition after the two lower branches can be constructed by permuting the three dimensions, e.g., $(n_1, n_2, n_3) \rightarrow (n_2, n_3, n_1)$ and $(\omega_1, \omega_2, \omega_3) \rightarrow (\omega_2, \omega_3, \omega_1)$, from the corresponding channels in the top branch. This applies to both the sampling matrices and the filters used in the decomposition.

C. Subsequent Levels of Decomposition

Fig. 7 shows the block diagram of subsequent levels of decompositions on one of the three branches. After the 3-D hourglass filter ($P_{0,0}^{(0)}(\omega)$), we sequentially decompose the signal by two 2-D filter banks, with the first one, denoted as $\text{IRC}_{12}^{(l_2)}$, operating along the (n_1, n_2) -plane and the second one, $\text{IRC}_{13}^{(l_3)}$, along the (n_1, n_3) -plane.

The 2-D filter bank $\text{IRC}_{12}^{(l_2)}$, which stands for Iteratively Resampled Checkerboard filter bank, has a binary-tree structure with l_2 (≥ 0) levels of decomposition and, therefore, has 2^{l_2} different output branches. The second filter bank $\text{IRC}_{13}^{(l_3)}$ has the same construction as $\text{IRC}_{12}^{(l_2)}$, but operates along a different signal plane, i.e., $(n_1, n_2) \rightarrow (n_1, n_3)$, and with a different decomposition depth, i.e., $l_2 \rightarrow l_3$. Note that we attach an $\text{IRC}_{13}^{(l_3)}$ to every output channel of $\text{IRC}_{12}^{(l_2)}$, so we have a total of $2^{l_2+l_3}$ output channels in Fig. 7. Postponing the detailed construction of the IRC filter banks to the next section, here, we first study the conditions the IRCs must satisfy so that the overall filter bank in Fig. 7, indeed, achieves the desired pyramid-shaped frequency decomposition given in Fig. 5(c).

Theorem 1: We index the $2^{l_2+l_3}$ output channels in NDFB from top to bottom with a pair of integers (i, j) , where $0 \leq i < 2^{l_2}$ and $0 \leq j < 2^{l_3}$. Meanwhile, we use $\tilde{P}_{i,j}^{(l_2,l_3)}(\omega_1, \omega_2, \omega_3)$ to denote the ideal equivalent subband filter at the (i, j) th channel, and $P_{i,j}^{(l_2,l_3)}(\omega_1, \omega_2, \omega_3)$ for the ideal pyramid-shaped filters. We have

$$\tilde{P}_{i,j}^{(l_2,l_3)}(\omega_1, \omega_2, \omega_3) = P_{i,j}^{(l_2,l_3)}(\omega_1, \omega_2, \omega_3), \quad \text{for all } i, j \text{ and } l_2, l_3$$

if the IRC filter banks satisfy the following two conditions.

- 1) **Equivalent sampling matrices.** Denote $\mathbf{M}_k^{(l_2)}$ as the overall sampling matrix for the k th ($0 \leq k < 2^{l_2}$) channel in the l_2 -level filter bank $\text{IRC}_{12}^{(l_2)}$, then

$$\mathbf{M}_k^{(l_2)} = \begin{pmatrix} 1 & 0 \\ 0 & 2^{l_2} \end{pmatrix}. \quad (3)$$

- 2) **Wedge refinement.** Denote $F_k^{(l_2)}(\omega_1, \omega_2)$ as the equivalent filter for the k th ($0 \leq k < 2^{l_2}$) channel in $\text{IRC}_{12}^{(l_2)}$, then

$$W_k^{(l_2)}(\omega_1, \omega_2) = W_0^{(0)}(\omega_1, \omega_2) \cdot F_k^{(l_2)}(\omega_1, \omega_2). \quad (4)$$

for all $l_2 \geq 0$.

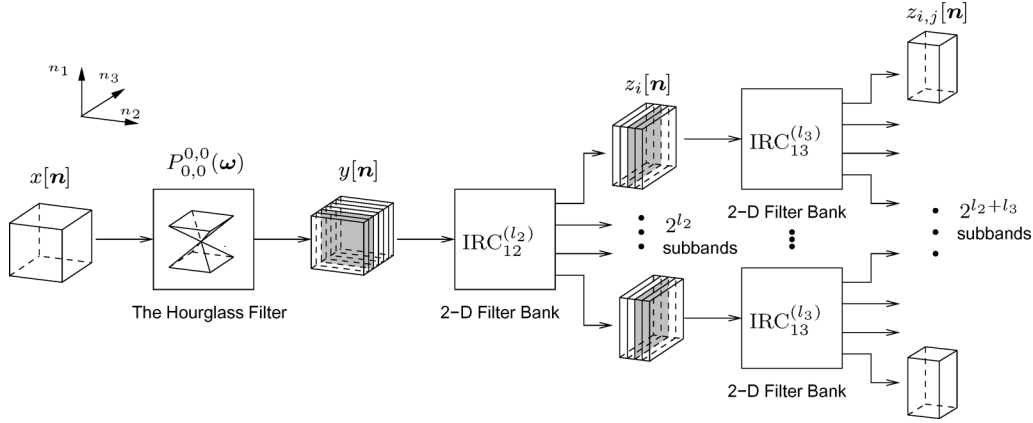


Fig. 7. One branch of the proposed filter bank structure of the NDFB in 3-D. The input signal $x[\mathbf{n}]$ first goes through the 3-D hourglass filter $P_{0,0}^{(0,0)}(\omega)$. The output $y[\mathbf{n}]$ is then fed into a 2-D filter bank, denoted by $\text{IRC}_{12}^{(l_2)}$, which operates on the (n_1, n_2) planes. The l_2 -level tree-structured filter bank $\text{IRC}_{12}^{(l_2)}$ produces 2^{l_2} output subbands, denoted as $z_i[\mathbf{n}]$ for $0 \leq i < 2^{l_2}$. Each output is then fed into another 2-D filter bank $\text{IRC}_{13}^{(l_3)}$ operating on the (n_1, n_3) planes. In the end, we get $2^{l_2+l_3}$ outputs, represented by $z_{i,j}[\mathbf{n}]$ for $0 \leq i < 2^{l_2}$ and $0 \leq j < 2^{l_3}$.

Proof: The flow of operations in each channel of Fig. 7 can be illustrated as follows:

$$\begin{aligned} \rightarrow P_{0,0}^{(0,0)}(\omega_1, \omega_2, \omega_3) &\rightarrow \underbrace{F_i^{(l_2)}(\omega_1, \omega_2)}_{\text{on } (n_1, n_2) \text{ planes}} \rightarrow (\downarrow \mathbf{M}_i^{(l_2)}) \rightarrow \bullet \\ &\bullet \rightarrow \underbrace{F_j^{(l_3)}(\omega_1, \omega_3)}_{\text{on } (n_1, n_3) \text{ planes}} \rightarrow (\downarrow \mathbf{M}_j^{(l_3)}) \rightarrow \bullet \end{aligned}$$

In calculating the overall equivalent filter of that channel, the filter $F_j^{(l_3)}(\omega_1, \omega_3)$ from the second filter bank $\text{IRC}_{13}^{(l_3)}$ will not be affected by the downsampling operation in the first filter bank $\text{IRC}_{12}^{(l_2)}$, because by the sampling matrices condition in (3), the sampling matrix $\mathbf{M}_i^{(l_2)}$ of $\text{IRC}_{12}^{(l_2)}$ is a diagonal matrix with its first element (along the n_1 dimension) being one. Therefore, we have

$$\begin{aligned} \tilde{P}_{i,j}^{(l_2, l_3)}(\omega_1, \omega_2, \omega_3) &= P_{0,0}^{(0,0)}(\omega_1, \omega_2, \omega_3) \\ &\quad \cdot F_i^{(l_2)}(\omega_1, \omega_2) \cdot F_j^{(l_3)}(\omega_1, \omega_3) \end{aligned} \quad (5)$$

for all $l_2, l_3 \geq 0$ and $0 \leq i < 2^{l_2}$, $0 \leq j < 2^{l_3}$. Using the equality in (1), we can decompose the hourglass filter in (5) as the product of two wedge filters

$$P_{0,0}^{(0,0)}(\omega_1, \omega_2, \omega_3) = W_0^{(0)}(\omega_1, \omega_2) \cdot W_0^{(0)}(\omega_1, \omega_3).$$

Now (5) can be rewritten as

$$\begin{aligned} \tilde{P}_{i,j}^{(l_2, l_3)}(\omega_1, \omega_2, \omega_3) &= \left(W_0^{(0)}(\omega_1, \omega_2) \cdot F_i^{(l_2)}(\omega_1, \omega_2) \right) \\ &\quad \cdot \left(W_0^{(0)}(\omega_1, \omega_3) \cdot F_j^{(l_3)}(\omega_1, \omega_3) \right). \end{aligned} \quad (6)$$

It then follows from the wedge refinement condition given in (4) that

$$\begin{aligned} \tilde{P}_{i,j}^{(l_2, l_3)}(\omega_1, \omega_2, \omega_3) &= W_i^{(l_2)}(\omega_1, \omega_2) \cdot W_j^{(l_3)}(\omega_1, \omega_3) \\ &= P_{i,j}^{(l_2, l_3)}(\omega_1, \omega_2, \omega_3) \end{aligned}$$

where the second equality is due to (1) again. ■

D. Construction and Properties of the IRC Filter Banks

In this section, we present a detailed construction of the 2-D filter banks $\text{IRC}_{12}^{(l_2)}$, and show that they indeed satisfy the two conditions required in Theorem 1.

When $l_2 = 0$, the filter bank $\text{IRC}_{12}^{(0)}$ is simply the identity transform. We need to consider this degenerate case, since sometimes we just want to decompose the 3-D hourglass support along only one, e.g., the ω_3 , direction.

When $l_2 \geq 1$, the analysis part of the filter bank $\text{IRC}_{12}^{(l_2)}$ is constructed as an l_2 -level binary tree by recursively attaching a copy of the diagram contents enclosed by the dashed rectangle in Fig. 8(a) to every output channels from the previous level [see Fig. 8(b)]. Furthermore, we attach a resampling matrix $\mathbf{U}_k^{(l_2)}$ ($0 \leq k < 2^{l_2}$) to each of the 2^{l_2} output channels of the tree.

As illustrated in Fig. 8(a), the building block of the tree is a two-channel 2-D filter bank with a checkerboard-shaped frequency partition. Note that we need to attach two resampling operations, denoted as \mathbf{R}_0 and \mathbf{R}_1 , to channel 0 and channel 1, respectively. The sampling matrices in Fig. 8(a) and (b) are defined as

$$\mathbf{D}_2 = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}, \quad \mathbf{R}_0 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{R}_1 = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}$$

$$\text{and } \mathbf{U}_k^{(l_2)} = \mathbf{R}_1^{2^{l_2}-1-2k}.$$

We can index the channels of the $\text{IRC}_{12}^{(l_2)}$ from top to bottom with integers from 0 to $2^{l_2} - 1$. Associated with each channel indexed by k ($0 \leq k \leq 2^{l_2} - 1$) is a sequence of path types $(t_1, t_2, \dots, t_{l_2})$, where a type t_i is either 0 for the upper branch or 1 for the lower branch, as shown in Fig. 8(a) and (b). According to the expanding rule, $(t_1, t_2, \dots, t_{l_2})$ is the binary representation of k , i.e.,

$$k = \sum_{i=1}^{l_2} t_i 2^{l_2-i}, \quad t_i \in \{0, 1\}.$$

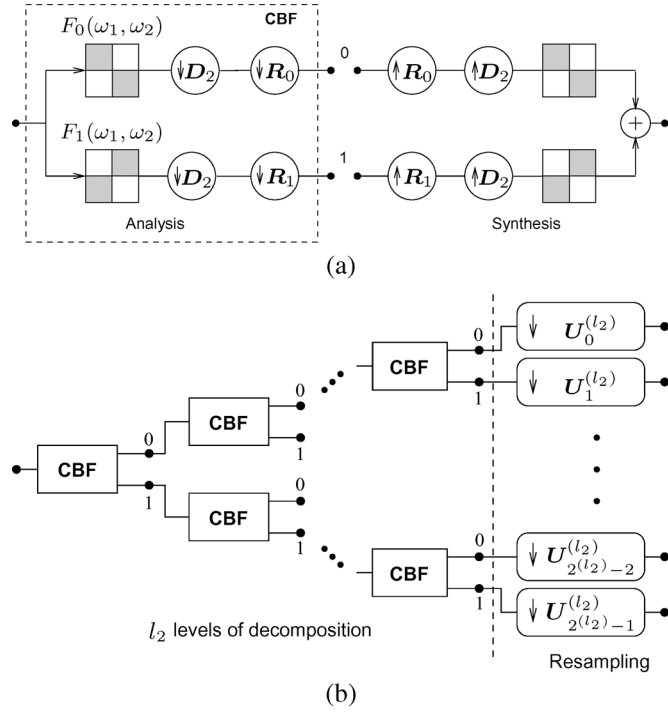


Fig. 8. (a) Two-channel 2-D checkerboard filter bank with resampling. The dark regions represent the ideal passband. (b) The filter bank $\text{IRC}_{12}^{(l_2)}$ is an l_2 -level tree-structured expansion of the checkerboard filter bank (denoted by CBF) given in (a), with resampling matrices $U_k^{(l_2)}$ attached at the end of each channel.

With this path type, the sequence of filtering and downsampling for channel k can be written as

$$\begin{aligned} \rightarrow [F_{t_1}] &\rightarrow (\downarrow D_2 R_{t_1}) \rightarrow [F_{t_2}] \rightarrow (\downarrow D_2 R_{t_2}) \rightarrow \dots \\ &\rightarrow [F_{t_{l_2}}] \rightarrow (\downarrow D_2 R_{t_{l_2}}) \rightarrow (\downarrow U_k^{(l_2)}) \end{aligned}$$

where F_0 and F_1 are the two component filters in the 2-D checkerboard filter bank as shown in Fig. 8(a).

From this, using the multirate identities recursively, we can transform the analysis side of the channel k ($0 \leq k < 2^{l_2}$) of the $\text{IRC}_{12}^{(l_2)}$ into a single filtering with the equivalent filter $F_k^{(l_2)}(\omega)$ followed by downsampling by the overall sampling matrix $M_k^{(l_2)}$. The following two propositions verify that the equivalent filters and sampling matrices indeed satisfy the conditions required in Theorem 1.

Proposition 1: The overall sampling matrix for the k th ($0 \leq k < 2^{l_2}$) channel in the l_2 -level filter bank $\text{IRC}_{12}^{(l_2)}$ is

$$M_k^{(l_2)} = \begin{pmatrix} 1 & 0 \\ 0 & 2^{l_2} \end{pmatrix}. \quad (7)$$

Proof: See Appendix A. ■

Remark: The sampling matrix $M_k^{(l_2)}$ has a determinant of 2^{l_2} , which makes the l_2 -level 2-D filter bank $\text{IRC}_{12}^{(l_2)}$ maximally decimated (nonredundant). However, the key is that all the downsampling operation is done on the n_2 -dimension, and this ensures the following filter bank $\text{IRC}_{13}^{(l_3)}$, which operates on the (n_1, n_3) dimensions, will not be affected by the downsampling. Actually, this is one of the distinctive features of our

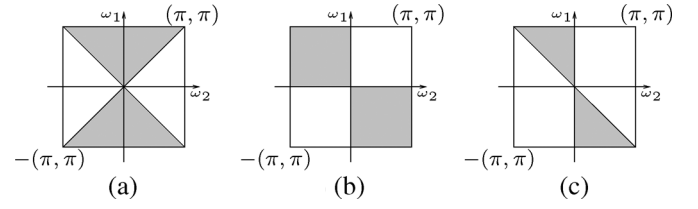


Fig. 9. (a) Ideal frequency support of $W_0^{(0)}(\omega_1, \omega_2)$. (b) The ideal checkerboard-shaped frequency support of $F_0(\omega_1, \omega_2)$. (c) Multiplying the supports in (a) and (b), we get the desired wedge-shaped frequency support of $W_0^{(1)}(\omega_1, \omega_2)$.

proposed filter bank in comparison to other 3-D constructions [10], [12].

In terms of the wedge refinement condition in (4), we want to verify that the wedge-shaped frequency support as shown in Fig. 5(a) can be achieved by applying the 2-D filter bank $\text{IRC}_{12}^{(l_2)}$. The special case when $l_2 = 1$ is illustrated in Fig. 9, where a wedge-shaped support $W_0^{(0)}(\omega_1, \omega_2)$ [see Fig. 9(a)] is divided by the checkerboard filter [see Fig. 9(b)] provided by $\text{IRC}_{12}^{(1)}$. The result is a “thinner” wedge support $W_0^{(1)}(\omega_1, \omega_2)$ shown in Fig. 9(c). We can show a general result holds for $l_2 \geq 0$, as stated in the following proposition.

Proposition 2: Assume the l_2 -level filter bank $\text{IRC}_{12}^{(l_2)}$ uses ideal filters. The equivalent filter $F_k^{(l_2)}(\omega_1, \omega_2)$ for the k th channel, $0 \leq k < 2^{l_2}$, satisfies the wedge refinement condition, i.e.,

$$W_k^{(l_2)}(\omega_1, \omega_2) = W_0^{(0)}(\omega_1, \omega_2) \cdot F_k^{(l_2)}(\omega_1, \omega_2)$$

for all $l_2 \geq 0$.

Proof: See Appendix B. ■

IV. NDFB IN ARBITRARY N -DIMENSIONAL ($N \geq 2$) CASES

In this section, we consider the NDFB construction in the general N -D cases for ($N \geq 2$). It turns out that the extension is surprisingly simple and requires no further filter bank design beyond the 3-D case.

A. N -D Directional Frequency Supports

In the general N -dimensional case, the ideal frequency supports of the NDFB are hypercube-based hyperpyramids. Analogous to the 3-D case, the first level of decomposition for NDFB is achieved by an N -channel undecimated filter bank, whose component filters have N -D “hourglass”-shaped frequency supports aligned with the $\omega_1, \dots, \omega_N$ axes, respectively. For the special case when $N = 2$, this is just a two-channel filter bank with a frequency partitioning as shown in Fig. 9(a).

From the geometric symmetry among the N branches, we only need to consider subsequent decompositions on the first branch, while the construction on the other branches can be inferred via a circular shift of variables. For instance, we can obtain the construction for the i th ($i \geq 2$) branch by changing the variables as follows: $(\omega_1, \dots, \omega_N) \rightarrow (\omega_i, \dots, \omega_N, \omega_1, \dots, \omega_{i-1})$ and $(n_1, \dots, n_N) \rightarrow (n_i, \dots, n_N, n_1, \dots, n_{i-1})$.

To describe the high-dimensional passband shapes, we also need to introduce some new notations. First, we evenly divide

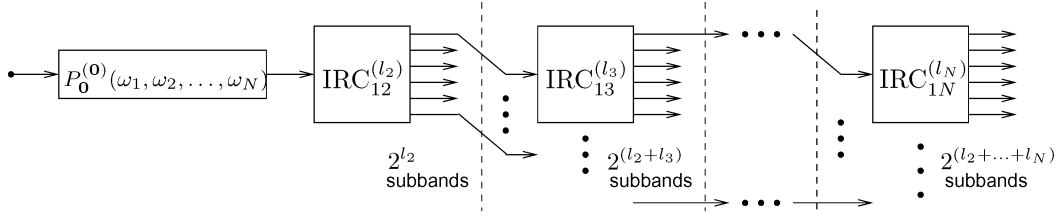


Fig. 10. Analysis part of the NDFB in the N -dimensional case. Only the top branch is shown, while the filter banks for the lower $N - 1$ branches can be obtained by permuting the variables from the corresponding channels in the top branch. Note that only the first stage, filtering by $P_0^{(0)}(\omega_1, \omega_2, \dots, \omega_N)$, is in N -D; all subsequent stages are 2-D operations. More specifically, $\text{IRC}_{1i}^{(l_i)}$ operates on 2-D slices of the input N -D signal represented by the dimension pair (n_1, n_i) .

$[-\pi, \pi)$ into 2^l ($l \geq 0$) line segments, and let $B_k^{(l)}$ to denote the k th segment, i.e.,

$$B_k^{(l)} \stackrel{\text{def}}{=} \left[-\pi + \frac{2\pi}{2^l}k, -\pi + \frac{2\pi}{2^l}(k+1) \right), \quad \text{for } 0 \leq k < 2^l.$$

Hyperpyramids: We can represent those hypercube-based hyperpyramid frequency supports aligned with the ω_1 axis by the following sets:

$$\mathcal{P}_{\mathbf{k}}^{(l)} \stackrel{\text{def}}{=} \left\{ \boldsymbol{\omega} = \alpha(\pi, b_2, \dots, b_N)^T : |\alpha| \leq 1 \text{ and } b_i \in B_{k_i}^{(l_i)} \right\}$$

where $\mathbf{k} = (k_2, \dots, k_N)$, $\mathbf{l} = (l_2, \dots, l_N)$, and $0 \leq k_i < 2^{l_i}$ for $i = 2, \dots, N$. In the 3-D case (i.e., $N = 3$) and with $l_2 = l_3 = 2$, we can easily verify that the $\mathcal{P}_{\mathbf{k}}^{(l)}$ defined above are just the set theoretic representations of the 16 different pyramid-shaped support regions shown in Fig. 5(c).

The ideal pyramid filter can be defined as

$$P_{\mathbf{k}}^{(l)}(\boldsymbol{\omega}) \stackrel{\text{def}}{=} 1_{\mathcal{P}_{\mathbf{k}}^{(l)}}(\boldsymbol{\omega}), \quad \text{for } \boldsymbol{\omega} \in [-\pi, \pi)^N \quad (8)$$

where $1_{\mathcal{S}}(\cdot)$ is the indicator function on the set \mathcal{S} ; i.e., $1_{\mathcal{S}}(\boldsymbol{\omega}) = 1$ if $\boldsymbol{\omega} \in \mathcal{S}$ and $1_{\mathcal{S}}(\boldsymbol{\omega}) = 0$ if otherwise.

Wedges: The frequency support of the 2-D wedge filters operating along the (ω_1, ω_i) plane ($i = 2, \dots, N$) can be represented as

$$\mathcal{W}_{k_i}^{(l_i)} \stackrel{\text{def}}{=} \left\{ (\omega_1, \omega_i) = \alpha(\pi, b_i) : |\alpha| \leq 1 \text{ and } b_i \in B_{k_i}^{(l_i)} \right\} \quad (9)$$

and, hence, the corresponding ideal 2-D wedge filter can be defined as

$$W_{k_i}^{(l_i)}(\omega_1, \omega_i) \stackrel{\text{def}}{=} 1_{\mathcal{W}_{k_i}^{(l_i)}}(\omega_1, \omega_i), \quad \text{for } (\omega_1, \omega_i) \in [-\pi, \pi)^2.$$

For example, $W_{k_2}^{(2)}(\omega_1, \omega_2)$ (with $k_2 = 0, 1, 2, 3$) represent the four 2-D wedge filters whose supports are shown in Fig. 5(a).

As an N -D generalization of the equality given in (1), we have the following result, simply stating that the ideal N -D hyperpyramid filters can be factorized as the product of $N - 1$ ideal 2-D wedge filters.

Lemma 1: The hyperpyramid filters defined in (8) are second-order generalized separable. Specifically

$$P_{\mathbf{k}}^{(l)}(\omega_1, \dots, \omega_N) = \prod_{i=2}^N W_{k_i}^{(l_i)}(\omega_1, \omega_i)$$

for all $\mathbf{l} = (l_2, \dots, l_N)$, $\mathbf{k} = (k_2, \dots, k_N)$ with $0 \leq k_i < 2^{l_i}$.

Proof: See Appendix C. ■

B. The NDFB Construction in Higher Dimensions

We show in Fig. 10 the construction of the analysis part of the proposed NDFB in the N -D case. The synthesis part (not shown in the figure) is exactly symmetric. After the first level of decomposition, which is an N -D “hourglass”-shaped filter $P_0^{(0)}(\omega_1, \omega_2, \dots, \omega_N)$, the input signal is further decomposed by a series of 2-D iteratively resampled checkerboard filter banks $\text{IRC}_{1i}^{(l_i)}$ ($i = 2, 3, \dots, N$), where $\text{IRC}_{1i}^{(l_i)}$ operates on 2-D slices of the input signal represented by the dimension pair (n_1, n_i) . Note that, starting from the second level, we attach an IRC filter bank to each output channel from the previous level, and, hence, the entire filter bank shown in Fig. 10 has a total of $2^{(l_2+\dots+l_N)}$ output channels.

For completeness, we state the following theorem, showing that the N -D filter bank in Fig. 10 indeed achieves the hyperpyramid-shaped frequency decomposition as defined in (8).

Theorem 2: We index the $2^{(l_2+\dots+l_N)}$ output channels in Fig. 10 from top to bottom with a sequence of $N - 1$ integers $\mathbf{k} = (k_2, \dots, k_N)$, where $0 \leq k_i < 2^{l_i}$. Meanwhile, we use $\tilde{P}_{\mathbf{k}}^{(l)}(\boldsymbol{\omega})$ to denote the equivalent subband filter at the \mathbf{k} th channel. Assume all filters are ideal, then we have

$$\tilde{P}_{\mathbf{k}}^{(l)}(\boldsymbol{\omega}) = P_{\mathbf{k}}^{(l)}(\boldsymbol{\omega})$$

for all $\mathbf{l} = (l_2, \dots, l_N)$, $\mathbf{k} = (k_2, \dots, k_N)$ with $0 \leq k_i < 2^{l_i}$.

The proof is similar to the one for the 3-D case and is omitted here. Again, the key ideas are to use the generalized separability of the hyperpyramid support given in Lemma 1 and the two properties of the IRC filter banks, i.e., overall downsampling matrices (Proposition 1) and wedge refinement (Proposition 2).

V. FILTER DESIGN AND IMPLEMENTATION ISSUES

In previous discussions, we assume all the filters used are ideal. In this section, we will design real filters that approximate the desired frequency responses.

A. Designing the Hourglass Filter Bank

The first task is to design the undecimated filter bank with perfect reconstruction and the desired hourglass-shaped frequency decomposition. To simplify exposition, we will focus on the 3-D case, while the design can be extended to arbitrary N -D cases straightforwardly.

Generally speaking, designing three and higher dimensional perfect reconstruction filter banks using FIR filters is a very challenging task with few ready-to-use tools available. However, we can greatly simplify the design problem if we do not confine ourselves to using FIR filters and work in the frequency

domain instead. As suggested by several independent work [23]–[25], multidimensional filters designed in the frequency domain can achieve quite satisfactory performance.

In this paper, we also design our hourglass filter bank in the frequency domain. As shown in Fig. 6, we use $H_i(\omega_1, \omega_2, \omega_3)$ and $G_i(\omega_1, \omega_2, \omega_3)$ for $i = 1, 2, 3$ to represent the three analysis and synthesis filters in the hourglass filter bank, respectively. As the first step of simplification, we assume the three analysis filters are rotated versions of each other, i.e.,

$$H_2(\omega_1, \omega_2, \omega_3) = H_1(\omega_2, \omega_3, \omega_1)$$

and

$$H_3(\omega_1, \omega_2, \omega_3) = H_1(\omega_3, \omega_1, \omega_2).$$

The same constraint also applies to the synthesis filters. Meanwhile, if the filter bank implements a tight frame expansion, we need the synthesis filters to be the time-reversed versions of the corresponding analysis filters, i.e.,

$$G_i(\omega) = H_i(-\omega) = H_i(\omega)$$

for $i = 1, 2, 3$, where the second equality comes from the symmetry in the ideal frequency responses of $H_i(\omega)$. Combining the above two constraints, we get the condition for perfect reconstruction as

$$H_1^2(\omega_1, \omega_2, \omega_3) + H_1^2(\omega_2, \omega_3, \omega_1) + H_1^2(\omega_3, \omega_1, \omega_2) = 1. \quad (10)$$

Inspired by the work of Feilner *et al.* [24] on building 2-D quincunx wavelets, we propose

$$H_1(\omega_1, \omega_2, \omega_3) = \sqrt{\frac{K(\omega_1, \omega_2, \omega_3)^\lambda}{K(\omega_1, \omega_2, \omega_3)^\lambda + K(\omega_2, \omega_3, \omega_1)^\lambda + K(\omega_3, \omega_1, \omega_2)^\lambda}} \quad (11)$$

where λ is a positive even integer and $K(\omega_1, \omega_2, \omega_3)$ is a positive and 2π periodic function of ω_1, ω_2 and ω_3 . We can verify that the perfect reconstruction condition in (10) is satisfied by arbitrary choices of λ and $K(\omega_1, \omega_2, \omega_3)$. To control the filter frequency responses so that they approximate the desired hourglass shape, we let

$$K(\omega_1, \omega_2, \omega_3) = E(\omega_1, \omega_2) \cdot E(\omega_1, \omega_3),$$

where $E(\cdot, \cdot)$ is a bivariate 2π periodic function such that $E(\omega_1, \omega_2)$ is approximately one in the dark region in Fig. 11(a) and zero in the white region, with smooth transition regions between the two. There can be many ways in designing $E(\cdot, \cdot)$. In our experiment, we employ the windowing method proposed by Tay and Kingsbury [26], in which we truncate the ideal “sinc”-like 2-D signal corresponding to the fan-shaped support in Fig. 11(a) with a smooth Kaiser window and take $E(\omega_1, \omega_2)$ to be the Fourier transform of the truncated signal. The parameter λ in (11) can be used to adjust the sharpness of the frequency response. In our experiment, we find $\lambda = 4$ to be a suitable value.

Fig. 11(b) shows the isosurface of the frequency response of one analysis filter. We can see that the frequency response approximates the ideal hourglass shape fairly well. Note that the

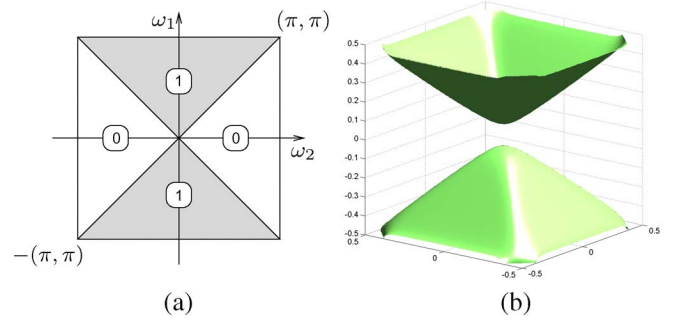


Fig. 11. (a) $E(\omega_1, \omega_2)$ approximately takes the value one in the dark region and the value zero in the white region. (b) The frequency response of one hourglass filter designed by the proposed frequency-domain method. The responses of the other two filters are rotational-symmetric to this one.

responses of the other two hourglass filters are rotated versions of this one. Finally, to extend this design to arbitrary N -dimensional cases, we can simply modify the parametrization to be

$$H_1(\omega) = \sqrt{\frac{\prod_{j=2}^N E(\omega_1, \omega_j)^\lambda}{\sum_{i=1}^N \prod_{j=1, j \neq i}^N E(\omega_i, \omega_j)^\lambda}}.$$

We find the hourglass filters designed here to produce satisfactory results in our numerical experiments. However, since these filters are not FIR or IIR, the filtering operations have to be implemented in the Fourier domain by first taking the FFT of the input signal and then multiplying it with the frequency values of the filters. In practice, this implementation usually requires a large memory space to store the entire multidimensional input signal¹ and causes long buffer delays which are undesirable for applications such as video processing. In [27], we report a different design of the hourglass filter banks using FIR filters, which allows the filtering operations to be carried out in the spatial domain with only partial input signal available in the memory.

B. Checkerboard Filter Bank

The other component in the NDFB construction is the 2-D checkerboard filter bank shown in Fig. 8(a). Unlike the 3-D hourglass filter bank, the design of the checkerboard filter bank has already been studied by several authors, including [26], [28]–[30]. In our experiment, we choose a design based on a two-stage ladder/lifting scheme [28], [30]. Attractive features of this design include good frequency selectivity using FIR filters, and efficient 1-D separable implementation in the polyphase domain. Due to space limitations, we omit here further description of the design and refer readers to [28], [30] for more details.

C. Implementation and Computational Cost

Because of the way we have specified the hourglass filters [see (11)], our implementation of the NDFB is entirely done in the Fourier domain based on FFT, including subsequent levels of decomposition. Although the checkerboard filter banks use FIR filters, they are still implemented in the Fourier domain to

¹It is possible to alleviate this memory requirement by employing an out-of-core implementation of the FFT, but at the expense of increased computation and I/O time.

TABLE I
RUNNING TIME (IN SECONDS) OF AN NDFB IMPLEMENTATION

| Data Size $N \times N \times N$ | # of Directional Subbands | | | |
|------------------------------------|---------------------------|--------------|---------------|---------------|
| | 3×1 | 3×4 | 3×16 | 3×64 |
| $128 \times 128 \times 128$ | 0.73 | 0.89 | 1.00 | 1.11 |
| $192 \times 192 \times 192$ | 2.88 | 3.78 | 4.03 | 4.33 |
| $256 \times 256 \times 256$ | 8.58 | 8.82 | 8.87 | 9.71 |

avoid the time-consuming FFT-IFFT operations in middle steps. To further improve computational efficiency, we take advantage of the fact that the input signal and filters are all real valued and, hence, only keep and operate on half of the FFT data.

Table I shows the performance of a C++ implementation of the NDFB on a computer with a 2.8 GHz processor and 1.5-GB memory. We perform the NDFB decomposition on 3-D signals of sizes $N \times N \times N$, where $N = 128, 192$ and 256 . The NDFB can also have different number of directional subbands, e.g., 3, 12, 48, 192 directions. We only show the running time for the forward transform (decomposition), while the time required by the inverse transform (reconstruction) is similar.

As we can see from the table, the FFT-based implementation is quite efficient. Meanwhile, for a fixed data size, the running time only increases in a roughly linear fashion when we quadruple the number of directional subbands, thanks to the tree-structured nature of the NDFB.

VI. SURFACELET TRANSFORM

A. Construction

As mentioned in Section I, the pyramid-shaped frequency partitioning makes NDFB a suitable tool in capturing surface singularities within multidimensional signals. In this section, we propose a multiscale version of the NDFB, called the *surfacelet transform*, to efficiently capture and represent local surface singularities with different sizes. This strategy is analogous to the contourlet construction [7], in which the original 2-D DFB is combined with a multiscale decomposition. However, an important distinction is that instead of using the Laplacian pyramid as in contourlets, we employ a new multiscale pyramid structure for the surfacelet transform, as shown in Fig. 12, which is conceptually similar to the one used in the steerable pyramid [1].

In Fig. 12, we use $L_i(\omega)$ ($i = 0, 1$) to represent the lowpass filters and $D_i(\omega)$ ($i = 0, 1$) to represent the highpass filters in the multiscale decomposition. $S(\omega)$ is an anti-aliasing filter used to cancel the aliasing caused by the upsampling operations. The NDFB is attached to the highpass branch at the finest scale and bandpass branches at coarser scales. To have more level of decomposition, we can recursively insert at point a_{n+1} a copy of the diagram contents enclosed by the dashed rectangle in the analysis part, and at point s_{n+1} a copy of the diagram contents enclosed by the dotted rectangle in the synthesis part.

In the new multiscale pyramid depicted in Fig. 12, the lowpass filter $L_0(\omega)$ in the first level is downsampled by a noninteger factor of 1.5 (upsampling by 2 followed by downsampling by 3) along each dimension. Although this fractional sampling factor makes the new pyramid slightly more redundant than the Laplacian pyramid (e.g., 1.34 versus 1.14 in redundancy ratios

in 3-D), we find the added redundancy to be very useful in reducing the frequency-domain aliasing of the NDFB, which is concentrated on the boundaries of the frequency cell $[-\pi, \pi)^N$ and mainly caused by the 2π periodicity of the frequency spectrum of discrete signals. Intuitively, the new multiscale pyramid achieves the task of eliminating aliasing components by only keeping the middle (alias-free) portion of the NDFB filter responses. Consequently, the constructed surfacelets are well-localized in both the spatial and frequency domain. Due to space limitation, we leave a detailed explanation for this issue to [31].

In our current implementation, we specify the lowpass filters $L_i(\omega)$ ($i = 0, 1$) in the frequency domain as $L_i(\omega) = d_i \cdot \prod_{n=1}^N L_i^{(1D)}(\omega_n)$, where $d_1 = 6^{N/2}$ and $d_2 = 2^{N/2}$; $L_i^{(1D)}(\omega_n)$ is a 1-D lowpass filter along the ω_n axis with passband frequency $\omega_{p,i}$, stopband frequency $\omega_{s,i}$, and a smooth transition band, defined as

$$L_i^{1D}(\omega) = \begin{cases} 1, & \text{for } |\omega| \leq \omega_{p,i} \\ \frac{1}{2} + \frac{1}{2} \cos\left(\frac{(|\omega| - \omega_{p,i})\pi}{\omega_{s,i} - \omega_{p,i}}\right), & \text{for } \omega_{p,i} < |\omega| < \omega_{s,i} \\ 0, & \text{for } \omega_{s,i} \leq |\omega| \leq \pi \end{cases} \quad (12)$$

for $|\omega| \leq \pi$ and $i = 0, 1$. Similarly, we also specify the anti-aliasing filter $S(\omega)$ as the separable product of 1-D lowpass filters having the same parametrized form as in (12), but with a different set of passband and stopband frequencies, denoted as $\omega_{p,A}$ and $\omega_{s,A}$, respectively.

The main advantage of designing the filters in the frequency domain is that we can let their frequency responses to be strictly zero beyond some cutoff frequencies. By choosing $\omega_{s,i}$ ($i = 1, 2$), $\omega_{p,A}$ and $\omega_{s,A}$ properly, we can ensure that the aliasing introduced by the upsampling and downsampling operations will be completely cancelled, and the perfect reconstruction condition for the multiscale pyramid can be simplified as

$$\frac{|L_i(\omega)|^2}{d_i^2} + |D_i(\omega)|^2 \equiv 1, \quad \text{for } i = 0, 1. \quad (13)$$

To satisfy the alias-free condition with an approximate octave-band decomposition, we choose the passband and stopband frequencies in (12) to be $\omega_{p,0} = \pi/3$, $\omega_{s,0} = 2\pi/3$, $\omega_{p,1} = \pi/4$, and $\omega_{s,1} = \pi/2$, and the frequency parameters of $S(\omega)$ to be $\omega_{p,A} = \pi/3$ and $\omega_{s,A} = 2\pi/3$. Once we have chosen the lowpass filters, the highpass filters $D_i(\omega)$ can be obtained from (13) to ensure perfect reconstruction.

B. Properties of the Surfacelet Transform

By combining the multiscale pyramid with the NDFB, the surfacelet transform has ideal passband supports as pairs of concentric cubes radiating out from the origin with different directions and scales. In the spatial domain, the surfacelet basis are localized surface patches with different normal directions and spatial locations. Table II summarizes the properties of the surfacelet transform, in comparison with two related systems that can also provide directional multiresolution signal decomposition in 3-D, including the real-valued dual-tree wavelet transform (DTWT) [5], [6] and the recently reported discrete implementation of the 3-D curvelet transform [25]. Note that there are two different versions of the 3-D curvelet transform, denoted

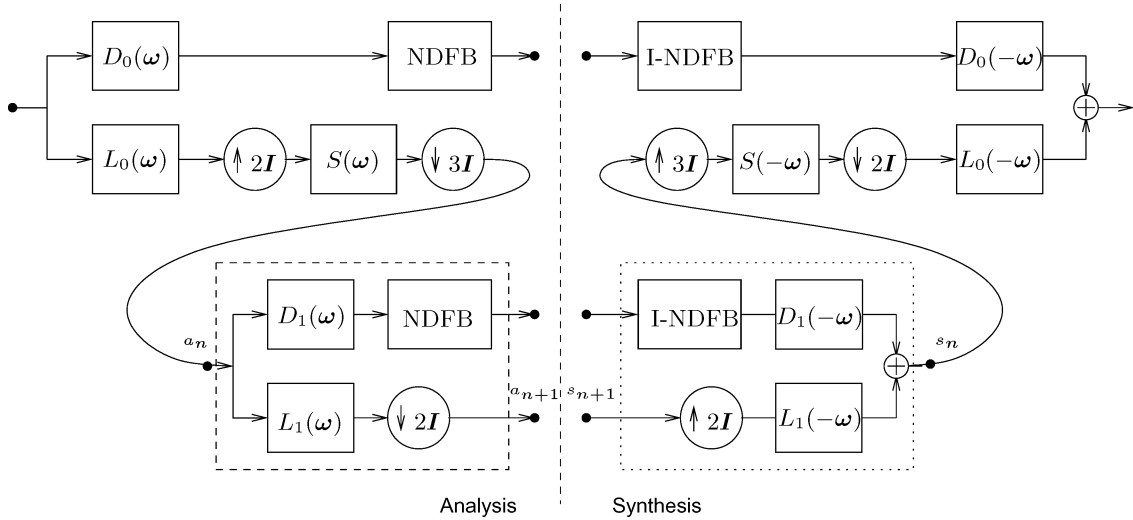


Fig. 12. Block diagram of the proposed surfacelet transform. The forward transform NDFB and its inverse INDFB are attached to the highpass subbands of the multiscale pyramid at each scale.

TABLE II
PROPERTIES OF THE SURFACELET TRANSFORM COMPARED WITH TWO OTHER DIRECTIONAL MULTIREOLUTION SYSTEMS

| | # of directional subbands at each scale | redundancy in 3-D | tree-structured construction | Implementation using FIR filters |
|--------------------|--|-------------------|------------------------------|----------------------------------|
| Dual-tree Wavelets | 28 | 4 | Yes | Yes |
| 3-D Curvelets-L | 3×2^l , except the finest scale | $4 \sim 5$ | No | No |
| 3-D Curvelets | 3×2^l , for all scales | ≈ 40 | No | No |
| Surfacelets | 3×2^l , for all scales | ≈ 4.02 | Yes | Possible with approximate PR |

in the table by curvelets and curvelets-L respectively, corresponding to whether wavelets or curvelets are used in the finest scale. Compared with the “full” version, curvelets-L has a substantially lower redundancy ratio and, hence, is more computationally feasible. However, since there is no directional decomposition at the finest scale, this version is more suitable for very bandlimited signals.

The surfacelet transform and the DTWT have similar redundancy ratios in 3-D. However, a potential advantage of the surfacelet transform is that its angular resolution can be refined (i.e. with more directional subbands) by invoking more levels of decomposition. In practice, we usually choose to have 192 or more directional subbands at finer scales, in contrast to the fixed 28 directional subbands provided in the DTWT.

The 3-D discrete curvelets and our proposed surfacelets aim at the same frequency partitioning, but the two transforms achieve this goal with two very different approaches. The curvelet implementation starts from defining a “mother” curvelet in the Fourier domain, whose scaled and sheared copies form a partition of unity. The curvelet coefficients are then obtained by multiplying the Fourier samples of the input signal with curvelet window functions at different scales and directions, followed by a spatial downsampling (implemented by frequency wrapping). Attractive features of this approach include its conceptual simplicity and direct connection with the continuous theory. However, an intrinsic problem is that all the downsampling is done in an alias-free fashion, requiring that the curvelet functions be strictly bandlimited in the frequency domain. This poses an inevitable trade-off between redundancy

and spatial localization of the curvelets, as lower redundancy requires higher subsampling ratio, which in turn implies narrower transition band in frequency supports, and from the uncertainty principle, slower decay of the resulting curvelets in the spatial domain.

In contrast, the NDFB, as the key component of the surfacelet transform, has a tree-structured filter bank construction, in which aliasing is allowed to exist and will be cancelled by carefully designed filters (see Section V). Consequently, the surfacelet transform is much less redundant than 3-D curvelets² (see Table II). Meanwhile, in NDFB we can use filters with fast spatial decay (and, hence, with more gently spread-out frequency support), without losing efficiency in terms of redundancy, since we do not need the filters to be strictly bandlimited. Furthermore, as a potential advantage of the surfacelet construction, it might be feasible to implement the multiscale pyramid shown in Fig. 12 using FIR filters, with the perfect reconstruction condition in (13) approximately satisfied. We leave the FIR implementation of the multiscale pyramid to a future work.

VII. EXPERIMENTAL RESULTS

In this section, we present some experiments with the proposed NDFB and the surfacelet transform. All experiments use the hourglass and checkerboard filters described in Section V and the multiscale filters specified in Section VI.

²When both use the same frequency partitioning at the finest scale.

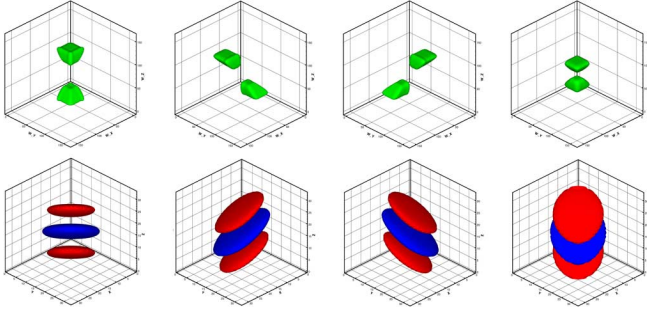


Fig. 13. Top row: Isosurfaces of four surfacelet basis images in the frequency domain. Bottom row: Isosurfaces of the same surfacelets, but in the spatial domain.

A. Surfacelets Basis Images

We show four surfacelets in the frequency domain in the top row of Fig. 13. The isosurface value is chosen to be half of the largest value in each basis image. As expected, the frequency support of surfacelets are bandpass concentric cubes tiling the frequency space with different directions.

Next, we show in the bottom row of Fig. 13 the same surfacelets, but in the spatial domain. In each basis image, the blue (or dark) colored isosurface is extracted at half of the most positive value and the red (or light) colored isosurface at half of the most negative value. We can see that the surfacelets in the spatial domain are localized surface patches, smooth along the tangent planes and oscillatory along the normal directions (equal to the directions of the corresponding frequency supports).

B. Three-Dimensional Zone Plate

We apply the proposed surfacelet transform to a 3-D zone plate image, which is generated by the formula $x[\mathbf{n}] = \cos(0.02 \times r^2)$, with $r = \sqrt{n_1^2 + n_2^2 + n_3^2}$ being the distance to the image center. As shown in Fig. 14(a), the zone plate $x[\mathbf{n}]$ represents a snapshot of a spherical wave at a certain time. Locally around a point $\mathbf{n}_0 = r \cdot \mathbf{u}$ with unit direction \mathbf{u} and radius r , the zone plate image can be approximated by a plane wave function that remains constant on all planes orthogonal to \mathbf{u} , and oscillates along \mathbf{u} at a frequency linearly proportional to r . In the frequency domain, the spectral support of that plane wave is localized around $\omega_0 = c \cdot r \cdot \mathbf{u}$ for some constant c . Comparing \mathbf{n}_0 with ω_0 , we can see that the spatial locations of the zone plate image directly correspond with its spectral locations.

In the experiment, we first decompose the 3-D image using the surfacelet transform, and then reconstruct the image by passing through only one of the subbands to the synthesis filter bank. As shown in Fig. 14(b)–(d), different subbands of the surfacelet transform can separate and capture the localized information of the 3-D zone plate with different directions and frequencies.

C. Video Denoising

Video can be seen as a special type of 3-D signals, with two spatial dimensions and one temporal dimension. Denoising video signals using velocity selective 3-D transforms has been studied by Selesnick *et al.* in [32]. In this experiment, we use

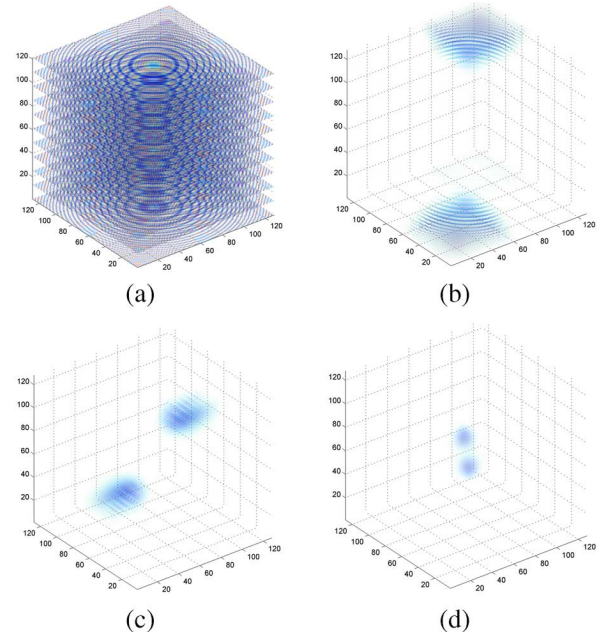


Fig. 14. (a) Original 3-D zone plate image; (b)–(d) the reconstructed images from three different subbands of the surfacelet transform at different scales and directions.

the proposed surfacelet transform to remove the zero mean white Gaussian noises added to several video signals. For benchmark, we also show the denoising performance of the 3-D undecimated separable wavelet transform (UDWT), the real-valued dual-tree wavelet transform (DTWT), and 3-D Curvelets-L. We use several standard SIF-sized test sequences for video coding and processing, including “Mobile,” “Coast-guard,” and “Tempete,” all of which can be obtained from <http://www.cipr.rpi.edu>.

In the test, the original sequences are truncated to the size of $192 \times 192 \times 192$. We use four levels of decomposition for all transforms. For the surfacelet transform, the number of directional subbands for each scale, from fine to coarse, is set to be 192, 192, 48, and 12. For the UDWT, we use the “symlet” of length 16. For a fair comparison, we employ the hard thresholding denoising method for all 4 transforms, by truncating the transform coefficients of the noisy sequences at the i th subband with a threshold T_i ($i = 1, 2, \dots$). Although not being the best denoising algorithm available, this simple hard thresholding scheme can often be a good indication of the potential of different transforms.

We choose the threshold $T_i = 3 \cdot E_i \cdot \hat{\sigma}_n$, where E_i is the precomputed L_2 norm of the equivalent filter at the i th subband, and $\hat{\sigma}_n$ is the estimated standard deviation of the input noise, obtained through a robust median estimator [33].

Table III shows the PSNR (in dB) of the denoised test sequences by using different transforms. To help interpret the results, we also list the redundancy ratio of each transform in the table, since the redundancy of a transform indicates its computational and memory efficiency, which is an important practical issue to consider in the context of 3-D signal processing.

It is not surprising that Curvelets-L does not perform as well as other transforms at low noise levels ($\sigma_n = 30$ or 40),

TABLE III
PSNR VALUES OF THE DENOISED SEQUENCES OF CURVELETS-L, UDWT, DTWT, AND THE PROPOSED SURFACELET TRANSFORM UNDER DIFFERENT NOISE LEVELS σ_n . THE REDUNDANCY OF EACH TRANSFORM IS SHOWN IN PARENTHESES

| σ_n | Mobile | | | Coastguard | | | Tempete | | |
|--------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | 30 | 40 | 50 | 30 | 40 | 50 | 30 | 40 | 50 |
| Curvelets-L (4.12) | 23.54 | 23.19 | 22.86 | 25.05 | 24.64 | 24.29 | 26.32 | 25.91 | 25.54 |
| UDWT (29) | 24.02 | 22.99 | 22.23 | 25.95 | 24.95 | 24.20 | 27.13 | 26.03 | 25.22 |
| DTWT (4) | 24.56 | 23.43 | 22.58 | 26.06 | 25.01 | 24.22 | 27.18 | 26.05 | 25.23 |
| Surfacelets (4.02) | 25.86 | 24.72 | 23.88 | 26.82 | 25.87 | 25.15 | 27.94 | 26.94 | 26.22 |

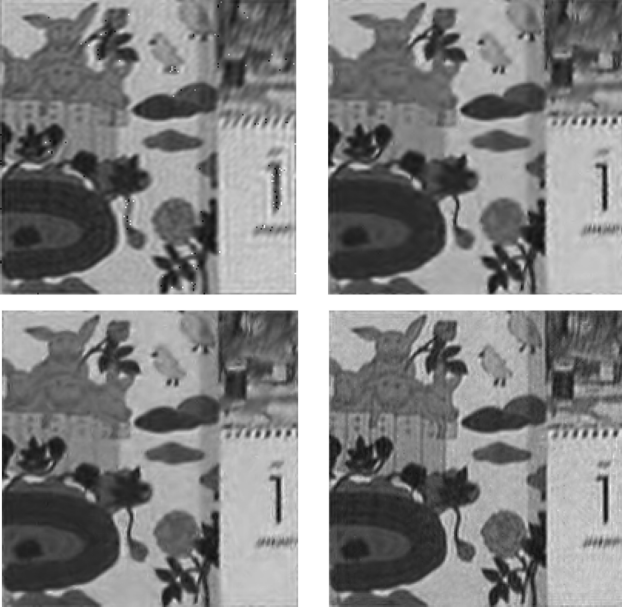


Fig. 15. Denoised frames from the “Mobile” sequence. From left to right: Curvelets-L (25.02 dB), UDWT (25.80 dB), DTWT (26.45 dB), and Surfacelets (28.29 dB). Shown in parentheses are the PSNR values calculated on each frame.

mainly because of its lack of directionality in the finest scale. We should mention that the “full” version of the 3-D curvelet transform (i.e., the one with directional decomposition at all scales) is expected to give better performance. However, its large redundancy (around 40) makes it beyond the capability of our testing platform, and prevents us from including it in the current experiment.

Among UDWT, DTWT and Surfacelets, we can see that Surfacelets outperforms the other two by a large margin (from 0.76 dB to 1.30 dB) for the tested sequences. This suggests the potential of the proposed surfacelet transform. Fig. 15 shows one frame from the denoised “Mobile” sequence by using different transforms. We can see that image details are best preserved by the surfacelet transform. This difference in denoising quality is much more conspicuous when viewing the video sequences.

VIII. CONCLUSION AND DISCUSSIONS

In this paper, we proposed a new family of directional filter banks for arbitrary N -dimensional signals. Compared with other related systems, the proposed NDFB is built upon an efficient tree-structured construction, which leads to a low redundancy ratio and refinable angular resolution. By combining the NDFB with a new multiscale pyramid, we constructed the surfacelet transform, which has potentials in efficiently

capturing and representing surface-like singularities in multidimensional signals. We envision that the proposed NDFB and surfacelet transform would find applications in various areas that involve the processing of multidimensional volumetric data, including video processing, seismic image processing, and medical image analysis.

There are several interesting directions worth pursuing. A constraint in the current construction is that the multiscale pyramid in the surfacelet transform has to be implemented in the Fourier domain. Therefore, a direction for future work is to find FIR implementations of the multiscale pyramid, perhaps with near perfect reconstruction. Another issue is that the current construction of NDFB is N -times redundant in the N -D case. Exploring new ways to further reduce this redundancy would be of practical importance, especially for higher dimensional cases. A software toolbox that implements the NDFB and surfacelet transform is freely available for download from the MATLAB Central (<http://www.mathworks.com/matlabcentral>).

APPENDIX

A. Proof of Proposition 1

We can write the equivalent downsampling matrix as

$$\begin{aligned} \mathbf{M}_k^{(l_2)} &= \left(\prod_{i=1}^{l_2} (\mathbf{D}_2 \cdot \mathbf{R}_{t_i}) \right) \cdot \mathbf{U}_k^{(l_2)} \\ &= \left(\prod_{i=1}^{l_2} (\mathbf{D}_2 \cdot \mathbf{R}_{t_i}) \right) \cdot \mathbf{R}_1^{2^{l_2}-1-2k}. \end{aligned} \quad (14)$$

Since $\mathbf{R}_1 = \mathbf{R}_0^{-1}$, we can write $\mathbf{R}_{t_i} = \mathbf{R}_0^{-2t_i+1}$ for $t_i = 0, 1$; thus, from (14)

$$\begin{aligned} \mathbf{M}_k^{(l_2)} &= \left(\prod_{i=1}^{l_2} (\mathbf{D}_2 \cdot \mathbf{R}_0^{-2t_i+1}) \right) \cdot \mathbf{R}_1^{2^{l_2}-1-2k} \\ &= \mathbf{D}_2^{l_2} \cdot \mathbf{R}_0^{s^{(l_2)}} \cdot \mathbf{R}_1^{2^{l_2}-1-2k} \\ &\quad \text{with } s^{(l_2)} = \sum_{i=1}^{l_2} (-2t_i + 1)2^{l_2-i} \end{aligned} \quad (15)$$

where the second equality is obtained by interchanging the positions of $\mathbf{R}_0^{-2t_i+1}$ and \mathbf{D}_2 in (15) recursively and using the fact that $\mathbf{R}_0^n \cdot \mathbf{D}_2 = \mathbf{D}_2 \cdot \mathbf{R}_0^{2n}$ for all $n \in \mathbb{Z}$. Recall that the channel index $k = \sum_{i=1}^{l_2} t_i 2^{l_2-i}$. After some simple manipulation, we can get

$$s^{(l_2)} = \sum_{i=1}^{l_2} (-2t_i + 1)2^{l_2-i} = 2^{l_2} - 1 - 2k$$

and, therefore, $\mathbf{M}_k^{(l_2)} = \mathbf{D}_2^{l_2} \cdot \mathbf{R}_0^{2^{l_2}-1-2k} \cdot \mathbf{R}_1^{2^{l_2}-1-2k} = \mathbf{D}_2^{l_2}$.

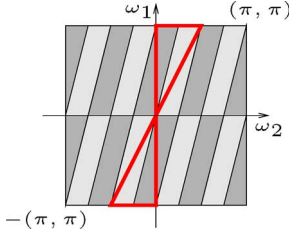


Fig. 16. Illustration of how the wedge support from the previous level can be decomposed into two finer slices by the resampled checkerboard filters. The region inside the thick lines represent one of the wedge supports from an l_2 -level decomposition. The dark gray and light gray regions inside the thin lines are copies of the basis spectrum of the upsampled checkerboard filters $F_0((Q_{l_2})^T \omega)$ and $F_1((Q_{l_2})^T \omega)$, respectively. The resampled (squeezed and sheared) checkerboard filters divide the original wedge right in the middle and produce two “thinner” wedges.

B. Proof of Proposition 2

We first write the equivalent filter as

$$F_k^{(l_2)}(\omega) = F_{t_1}(\omega) \prod_{n=2}^{l_2} F_{t_n}((Q_{n-1})^T \omega) \quad (16)$$

where $\omega = (\omega_1, \omega_2)^T$, and the matrix Q_{n-1} ($n = 2, \dots, l_2$) is understood as the partial product of the overall sampling matrix in (14), i.e.,

$$Q_{n-1} = \prod_{i=1}^{n-1} (D_2 \cdot R_{t_i}).$$

We will prove the proposition by induction. The cases for $l_2 = 0$ and $l_2 = 1$ can be easily verified pictorially, as in Fig. 9. Now suppose that the equivalent filters $F_{k'}^{(l_2)}(\omega)$, $0 \leq k' < 2^{l_2}$, for an l_2 -level IRC $_{12}^{(l_2)}$ satisfy the desired wedge refinement condition. The (l_2+1) -level filter bank IRC $_{12}^{(l_2+1)}$ is obtained by appending a resampled checkerboard filter bank to each channel of the IRC $_{12}^{(l_2)}$ (excluding the resampling matrices $U_{k'}^{l_2}$). We consider the equivalent filter $F_k^{(l_2+1)}(\omega)$ of the k th channel with $k = 2k' + t_{l_2+1}$, where $t_{l_2+1} = 0$ or 1 . From (16), we have

$$\begin{aligned} F_k^{(l_2+1)}(\omega) &= F_{t_1}(\omega) \prod_{n=2}^{l_2+1} F_{t_n}((Q_{n-1})^T \omega) \\ &= \left(F_{t_1}(\omega) \prod_{n=2}^{l_2} F_{t_n}((Q_{n-1})^T \omega) \right) F_{t_{l_2+1}}((Q_{l_2})^T \omega) \\ &= F_{k'}^{(l_2)}(\omega) \cdot F_{t_{l_2+1}}((Q_{l_2})^T \omega). \end{aligned} \quad (17)$$

Now by induction, we can write

$$\begin{aligned} W_0^{(0)}(\omega) \cdot F_k^{(l_2+1)}(\omega) &= W_0^{(0)}(\omega) \cdot F_{k'}^{(l_2)}(\omega) \cdot F_{t_{l_2+1}}((Q_{l_2})^T \omega) \\ &= W_{k'}^{(l_2)}(\omega) \cdot F_{t_{l_2+1}}((Q_{l_2})^T \omega). \end{aligned} \quad (18)$$

Using the expression for Q_{l_2} ($= D_2^{l_2} \cdot R_0^{2^{l_2}-1-2k}$) in Appendix A, it can be verified that in (18), upsampling by Q_{l_2} effectively squeezes and shears the basic spectrum of the checkerboard filter banks $F_{t_{l_2+1}}(\omega)$ [which were shown in Fig. 8(a)], so that the copy that overlaps with the wedge support $W_{k'}^{(l_2)}(\omega)$ exactly divides the original wedge in the middle. As illustrated in Fig. 16, the outcome are two “thinner” wedge supports that correspond to $W_{2k'}^{(l_2+1)}(\omega)$ and $W_{2k'+1}^{(l_2+1)}(\omega)$, respectively.

Now the right side of (18) can be expressed as $W_{2k'+t_{l_2+1}}^{(l_2+1)}(\omega) = W_k^{(l_2+1)}(\omega)$, and, therefore, we have verified the wedge refinement condition for the (l_2+1) -level case, i.e.,

$$W_0^{(0)}(\omega) \cdot F_k^{(l_2+1)}(\omega) = W_k^{(l_2+1)}(\omega).$$

C. Proof of Lemma 1

Taking the intersection of $N-1$ wedge supports, we have

$$\begin{aligned} &\bigcap_{i=2}^N W_{k_i}^{(l_i)} \\ &= \bigcap_{i=2}^N \left\{ (\omega_1, \omega_i) = \alpha_i(\pi, b_i) : |\alpha_i| \leq 1 \text{ and } b_i \in B_{k_i}^{(l_i)} \right\} \\ &= \bigcap_{i=2}^N \bigcup_{|\alpha_i| \leq 1} \alpha_i \cdot \left\{ (\omega_1, \omega_i) = (\pi, b_i) : b_i \in B_{k_i}^{(l_i)} \right\} \\ &= \bigcup_{|\alpha_2| \leq 1} \dots \bigcup_{|\alpha_N| \leq 1} \bigcap_{i=2}^N \alpha_i \cdot \left\{ (\omega_1, \omega_i) = (\pi, b_i) : b_i \in B_{k_i}^{(l_i)} \right\} \\ &= \bigcup_{|\alpha| \leq 1} \bigcap_{i=2}^N \alpha \cdot \left\{ (\omega_1, \omega_i) = (\pi, b_i) : b_i \in B_{k_i}^{(l_i)} \right\} \\ &= \bigcup_{|\alpha| \leq 1} \alpha \cdot \left\{ \omega = (\pi, b_2, \dots, b_N)^T : b_i \in B_{k_i}^{(l_i)} \right\} = \mathcal{P}_k^{(l)}. \end{aligned}$$

Since the ideal hyperpyramid and wedge filters are just indicator functions defined on their corresponding support sets, we have $P_k^{(l)}(\omega) = 1_{\mathcal{P}_k^{(l)}}(\omega) = \prod_{i=2}^N 1_{W_{k_i}^{(l_i)}}(\omega_1, \omega_i) = \prod_{i=2}^N W_{k_i}^{(l_i)}(\omega_1, \omega_i)$.

ACKNOWLEDGMENT

The authors would like to thank L. Demanet and L. Ying for helpful discussions on their 3-D curvelet software, and I. Selesnick for providing his 3-D dual-tree wavelet software.

REFERENCES

- [1] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, “Shiftable multiscale transforms,” *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 587–607, Mar. 1992.
- [2] R. H. Bamberger and M. J. T. Smith, “A filter bank for the directional decomposition of images: Theory and design,” *IEEE Trans. Signal Process.*, vol. 40, no. 4, pp. 882–893, Apr. 1992.
- [3] J. P. Antoine, P. Carrette, R. Murenzi, and B. Piette, “Image analysis with two-dimensional continuous wavelet transform,” *Signal Process.*, vol. 31, pp. 241–272, 1993.
- [4] E. J. Candès and D. L. Donoho, “Curvelets—A suprisingly effective nonadaptive representation for objects with edges,” in *Curve and Surface Fitting*, A. Cohen, C. Rabut, and L. L. Schumaker, Eds. Nashville, TN: Vanderbilt Univ. Press, 1999.
- [5] N. Kingsbury, “Complex wavelets for shift invariant analysis and filtering of signals,” *J. Appl. Comput. Harmon. Anal.*, vol. 10, pp. 234–253, 2001.
- [6] I. W. Selesnick, “The double-density dual-tree DWT,” *IEEE Trans. Signal Process.*, vol. 52, no. 5, pp. 1304–1314, May 2004.
- [7] M. N. Do and M. Vetterli, “The contourlet transform: An efficient directional multiresolution image representation,” *IEEE Trans. Image Process.*, vol. 14, no. 12, pp. 2091–2106, Dec. 2005.
- [8] E. L. Pennec and S. Mallat, “Sparse geometric image representations with bandelets,” *IEEE Trans. Image Process.*, vol. 14, no. 4, pp. 423–438, Apr. 2005.

- [9] K. Guo, W.-Q. Lim, D. Labate, G. Weiss, and E. Wilson, "Wavelets with composite dilations and their MRA properties," *J. Appl. Comput. Harmon. Anal.*, vol. 20, pp. 231–249, 2006.
- [10] R. H. Bamberger, "New results on two and three dimensional directional filter banks," in *Proc. 27th Asilomar Conf. Signals, Systems, Computers*, 1993, vol. 2, pp. 1286–1290.
- [11] S. Park, M. J. T. Smith, and R. M. Mersereau, "Improved structures of maximally decimated directional filter banks for spatial image analysis," *IEEE Trans. Image Process.*, vol. 13, no. 11, pp. 1424–1431, Nov. 2004.
- [12] S. Park, "New directional filter banks and their applications in image processing," Ph.D. dissertation, Georgia Inst. Technol., Atlanta, 1999.
- [13] M. N. Do, "Directional multiresolution image representations" Ph.D. dissertation, Swiss Federal Inst. Technol., Lausanne, Dec. 2001 [Online]. Available: <http://www.ifp.uiuc.edu/minhdo/publications>
- [14] P. S. Hong and M. J. T. Smith, "An octave-band family of nonredundant directional filter banks," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, Orlando, FL, 2002, pp. 1165–1168.
- [15] Y. Lu and M. N. Do, "CRISP-contourlets: A critically-sampled directional multiresolution image representation," presented at the SPIE Conf. Wavelet Applications in Signal and Image Processing X, San Diego, CA, Aug. 2003.
- [16] T. T. Nguyen and S. Orantara, "Multiresolution direction filter banks: Theory, design and applications," *IEEE Trans. Signal Process.*, vol. 53, no. 10, pp. 3895–3905, Oct. 2005.
- [17] L. T. Bruton, "Three-dimensional cone filter banks," *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.*, vol. 50, no. 2, pp. 208–216, Feb. 2003.
- [18] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA: SIAM, 1992.
- [19] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [20] S. Mallat, *A Wavelet Tour of Signal Processing*. New York: Academic, 1998.
- [21] E. Viscito and J. P. Allebach, "The analysis and design of multidimensional FIR perfect reconstruction filter banks for arbitrary sampling lattices," *IEEE Trans. Circuits Syst.*, vol. 38, no. 1, pp. 29–41, Jan. 1991.
- [22] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [23] F. Nicolier, O. Laligant, and F. Truchetet, "Discrete wavelet transform implementation in Fourier domain for multidimensional signal," *J. Electron. Imag.*, vol. 11, no. 3, pp. 338–346, Jul. 2002.
- [24] M. Feilner, D. V. D. Ville, and M. Unser, "An orthogonal family of quincunx wavelets with continuously adjustable order," *IEEE Trans. Image Process.*, vol. 14, no. 4, pp. 499–510, Apr. 2005.
- [25] L. Ying, L. Demanet, and E. Candès, "3D discrete curvelet transform," presented at the SPIE Conf. Wavelet Applications in Signal and Image Processing XI, San Diego, CA, 2005.
- [26] D. B. H. Tay and N. Kingsbury, "Flexible design of multidimensional perfect reconstruction FIR 2-band filters using transformations of variables," *IEEE Trans. Image Process.*, vol. 2, no. 4, pp. 466–480, Oct. 1993.
- [27] Y. Lu and M. N. Do, "Multidimensional nonsubsampling hourglass filter banks: Geometry of passband support and filter design," presented at the 40th Asilomar Conf. Signals, Systems, and Computers, Pacific Grove, CA, Oct. 2006.
- [28] Y.-P. Lin and P. P. Vaidyanathan, "Theory and design of two-dimensional filter banks: A review," *Multidimensional Syst. Signal Process.*, vol. 7, pp. 263–330, 1996.
- [29] J. G. Rosiles, "Image and texture analysis using biorthogonal angular filter banks," Ph.D. dissertation, Georgia Inst. Technol., Atlanta, 2004.
- [30] Y. Lu and M. N. Do, "The finer directional wavelet transform," presented at the IEEE Int. Conf. Acoustics, Speech, Signal Processing, Philadelphia, PA, Mar. 2005.
- [31] Y. Lu and M. N. Do, "A new contourlet transform with sharp frequency localization," presented at the IEEE Int. Conf. Image Processing, Atlanta, GA, Oct. 2006.
- [32] I. W. Selesnick and K. Y. Li, "Video denoising using 2D and 3D dual-tree complex wavelet transforms," in *SPIE Conf. Wavelet Applications in Signal and Image Processing X*, San Diego, CA, Aug. 2003.
- [33] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, Dec. 1994.



Yue M. Lu (S'04) received the B.Eng and M.Eng degrees in electrical engineering, both with highest honors, from Shanghai Jiao Tong University, China, in 1999 and 2002. He is currently pursuing the Ph.D. degree in electrical engineering and the M.S. degree in mathematics at the University of Illinois at Urbana-Champaign, Urbana.

He was a Visiting Student at Microsoft Research Asia, Beijing, China, from March to July 2002, and a Visiting Researcher at Siemens Corporate Research, Princeton, NJ, from January to June 2006. His research interests include the theory, constructions, and applications of multiscale geometric representations for multidimensional signals; image and video processing; and medical image analysis for computer-aided diagnosis.

Mr. Lu received the Most Innovative Paper Award of ICIP in 2006 for his paper (with M. N. Do) on the construction of directional multiresolution image representations.



Minh N. Do (M'01) was born in Thanh Hoa, Vietnam, in 1974. He received the B.Eng. degree in computer engineering from the University of Canberra, Australia, in 1997, and the Dr.Sci. degree in communication systems from the Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland, in 2001.

Since 2002, he has been an Assistant Professor with the Department of Electrical and Computer Engineering and a Research Assistant Professor with the Coordinated Science Laboratory and the Beckman Institute, University of Illinois at Urbana-Champaign, Urbana. His research interests include wavelets, image and multidimensional signal processing, multiscale geometric analysis, and visual information representation.

Dr. Do received a Silver Medal from the 32nd International Mathematical Olympiad in 1991, a University Medal from the University of Canberra in 1997, the best doctoral thesis award from the Swiss Federal Institute of Technology Lausanne in 2001, and a CAREER award from the National Science Foundation in 2003.