The Nonsubsampled Contourlet Transform: Theory, Design, and Applications

Arthur L. da Cunha, Jianping Zhou, *Member, IEEE*, and Minh N. Do, *Member, IEEE*

Abstract—In this paper, we develop the nonsubsampled contourlet transform (NSCT) and study its applications. The construction proposed in this paper is based on a nonsubsampled pyramid structure and nonsubsampled directional filter banks. The result is a flexible multiscale, multidirection, and shift-invariant image decomposition that can be efficiently implemented via the *à trous* algorithm. At the core of the proposed scheme is the nonseparable two-channel nonsubsampled filter bank (NSFB). We exploit the less stringent design condition of the NSFB to design filters that lead to a NSCT with better frequency selectivity and regularity when compared to the contourlet transform. We propose a design framework based on the mapping approach, that allows for a fast implementation based on a lifting or ladder structure, and only uses one-dimensional filtering in some cases. In addition, our design ensures that the corresponding frame elements are regular, symmetric, and the frame is close to a tight one. We assess the performance of the NSCT in image denoising and enhancement applications. In both applications the NSCT compares favorably to other existing methods in the literature.

Index Terms—Contourlet transform, frames, image denoising, image enhancement, multidimensional filter banks, nonsubsampled filter banks.

I. INTRODUCTION

A NUMBER of image-processing tasks are efficiently carried out in the domain of an invertible linear transformation. For example, image compression and denoising are efficiently done in the wavelet transform domain [1], [2]. An effective transform captures the essence of a given signal or a family of signals with few basis functions. The set of basis functions completely characterizes the transform and this set can be redundant or not, depending on whether the basis functions are linear dependent. By allowing redundancy, it is possible to enrich the set of basis functions so that the representation is more efficient in capturing some signal behavior. In addition, redundant representations are generally more flexible and easier to design. In applications such as denoising, enhancement, and contour detection, a redundant representation can significantly outperform a nonredundant one.

Another important feature of a transform is its stability with respect to shifts of the input signal. The importance of the shift-invariance property in imaging applications dates back at least to Daugman [3] and was also advocated by Simoncelli *et al.* in [4]. An example that illustrates the importance of shift-invariance is image denoising by *thresholding* where the lack of shift-invariance causes pseudo-Gibbs phenomena around singularities [5]. Thus, most state-of-the-art wavelet denoising algorithms (see for example [6]–[8]) use an expansion with less shift sensitivity than the standard maximally decimated wavelet decomposition—the most common being the nonsubsampled wavelet transform (NSWT) computed with the *à trous* algorithm [9].[1]

In addition to shift-invariance, it has been recognized that an efficient image representation has to account for the geometrical structure pervasive in natural scenes. In this direction, several representation schemes have recently been proposed [10]–[15]. The contourlet transform [14] is a multidirectional and multiscale transform that is constructed by combining the Laplacian pyramid [16], [17] with the directional filter bank (DFB) proposed in [18]. The pyramidal filter bank structure of the contourlet transform has very little redundancy, which is important for compression applications. However, designing good filters for the contourlet transfom is a difficult task. In addition, due to downsamplers and upsamplers present in both the Laplacian pyramid and the DFB, the contourlet transform is not shift-invariant.

In this paper, we propose an overcomplete transform that we call the *nonsubsampled contourlet transform* (NSCT). Our main motivation is to construct a flexible and efficient transform targeting applications where redundancy is not a major issue (e.g., denoising). The NSCT is a fully shift-invariant, multiscale, and multidirection expansion that has a fast implementation. The proposed construction leads to a filter-design problem that to the best of our knowledge has not been addressed elsewhere. The design problem is much less constrained than that of contourlets. This enables us to design filters with better frequency selectivity thereby achieving better subband decomposition. Using the mapping approach we provide a framework for filter design that ensures good frequency localization in addition to having a fast implementation through ladders steps. The NSCT has proven to be very efficient in image denoising and image enhancement as we show in this paper.

A. L. da Cunha 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: arthur.cunha@ifp.uiuc.edu).

J. Zhou was with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA. He is now with Texas Instruments, Austin, TX 78714 USA (e-mail: jianping@ti.com).

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).

[1]Denoising by thresholding in the NSWT domain can also be realized by denoising multiple circular shifts of the signal with a critically sampled wavelet transform and then averaging the results. This has been termed *cycle spinning* after [5].
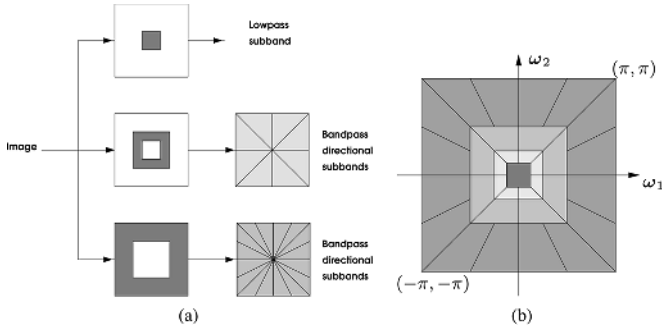
Fig. 1. Nonsubsampled contourlet transform. (a) NSFB structure that implements the NSCT. (b) Idealized frequency partitioning obtained with the proposed structure.

Fig. 2. Proposed nonsubsampled pyramid is a 2-D multiresolution expansion similar to the 1-D NSWT. (a) Three-stage pyramid decomposition. The lighter gray regions denote the aliasing caused by upsampling. (b) Subbands on the 2-D frequency plane.

The paper is structured as follows. In Section II, we describe the NSCT and its building blocks. We introduce a pyramid structure that ensures the multiscale feature of the NSCT and the directional filtering structure based on the DFB. The basic unit in our construction is the nonsubsampled filter bank (NSFB) which is discussed in Section II. In Section III, we study the issues associated with the NSFB design and implementation problems. Applications of the NSCT in image denoising and enhancement are discussed in Section IV. Concluding remarks are drawn in Section V.

Notation: Throughout this paper, a two-dimensional (2-D) filter is represented by its z-transform $H(\mathbf{z})$ where $\mathbf{z} = [z_1, z_2]^T$. Evaluated on the unit sphere, a filter is denoted by $H(e^{j\omega})$ where $e^{j\omega} = [e^{j\omega_1}, e^{j\omega_2}]^T$. If $\mathbf{m} = [m_1, m_2]^T$ is a 2-D vector, then $\mathbf{z}^{\mathbf{m}} = z_1^{m_1} z_2^{m_2}$ whereas if $\mathbf{M}$ is a 2×2 matrix, then $\mathbf{z}^{\mathbf{M}} = [\mathbf{z}^{\mathbf{m}_1}, \mathbf{z}^{\mathbf{m}_2}]$ with $\mathbf{m}_1, \mathbf{m}_2$ the columns of $\mathbf{M}$. In this paper we often deal with zero-phase 2-D filters. On the unit sphere, such filters can be written as polynomials in $\cos\omega = (\cos\omega_1, \cos\omega_2)^T$. We thus write $F(x_1, x_2)$ for a zero-phase filter in which x_1 and x_2 denote $\cos\omega_1$ and $\cos\omega_2$, respectively.

Abbreviations: A number of abbreviations are used throughout the paper.

　　NSCT: *Nonsubsampled Contourlet Transform.*
　　NSFB: *Nonsubsampled Filter Bank.*
　　NSDFB: *Nonsubsampled Directional Filter Bank.*
　　NSP: *Nonsubsampled Pyramid.*
　　NSWT: *Nonsubsampled 2-D Wavelet Transform.*
　　LAS: *Local Adaptive Shrinkage.*

II. NONSUBSAMPLED CONTOURLETS AND FILTER BANKS

A. Nonsubsampled Contourlet Transform

Fig. 1(a) displays an overview of the proposed NSCT. The structure consists in a bank of filters that splits the 2-D frequency plane in the subbands illustrated in Fig. 1(b). Our proposed transform can thus be divided into two shift-invariant parts: 1) a nonsubsampled pyramid structure that ensures the multiscale property and 2) a nonsubsampled DFB structure that gives directionality.

1) Nonsubsampled Pyramid (NSP): The multiscale property of the NSCT is obtained from a shift-invariant filtering structure that achieves a subband decomposition similar to that of the
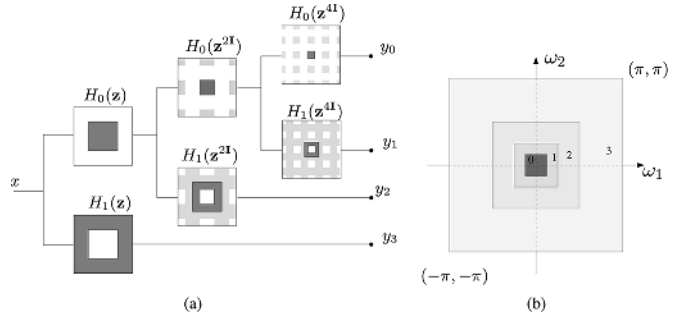
Laplacian pyramid. This is achieved by using two-channel nonsubsampled 2-D filter banks. Fig. 2 illustrates the proposed *nonsubsampled pyramid* (NSP) decomposition with $J = 3$ stages. Such expansion is conceptually similar to the one–dimensional (1-D) NSWT computed with the *à trous* algorithm [9] and has $J + 1$ redundancy, where J denotes the number of decomposition stages. The ideal passband support of the low-pass filter at the jth stage is the region $[-(\pi/2^j), (\pi/2^j)]^2$. Accordingly, the ideal support of the equivalent high-pass filter is the complement of the low-pass, i.e., the region $[(-\pi/2^{j-1}), (\pi/2^{j-1})]^2 \backslash [(-\pi/2^j), (\pi/2^j)]^2$. The filters for subsequent stages are obtained by upsampling the filters of the first stage. This gives the multiscale property without the need for additional filter design. The proposed structure is thus different from the separable NSWT. In particular, one bandpass image is produced at each stage resulting in $J + 1$ redundancy. By contrast, the NSWT produces three directional images at each stage, resulting in $3J + 1$ redundancy.

The 2-D pyramid proposed in [19] pp. 21 is obtained with a similar structure. Specifically, the NSFB of [19] is built from low-pass filter $H_0(\mathbf{z})$. One then sets $H_1(\mathbf{z}) = 1 - H_0(\mathbf{z})$, and the corresponding synthesis filters $G_0(\mathbf{z}) = G_1(\mathbf{z}) = 1$. A similar decomposition can be obtained by removing the downsamplers and upsamplers in the Laplacian pyramid and then upsampling the filters accordingly. Those perfect reconstruction systems can be seen as a particular case of our more general structure. The advantage of our construction is that it is general and as a result, better filters can be obtained. In particular, in our design $G_0(\mathbf{z})$ and $G_1(\mathbf{z})$ are low-pass and high-pass. Thus, they filter certain parts of the noise spectrum in the processed pyramid coefficients.

2) Nonsubsampled Directional Filter Bank (NSDFB): The directional filter bank of Bamberger and Smith [18] is constructed by combining critically-sampled two-channel fan filter banks and resampling operations. The result is a tree-structured filter bank that splits the 2-D frequency plane into directional wedges. A shift-invariant directional expansion is obtained with a nonsubsampled DFB (NSDFB). The NSDFB is constructed by eliminating the downsamplers and upsamplers in the DFB (see also [20]). This is done by switching off the downsamplers/upsamplers in each two-channel filter bank in the DFB tree structure and upsampling the filters accordingly. This results in a tree composed of two-channel NSFBs. Fig. 3 illustrates a
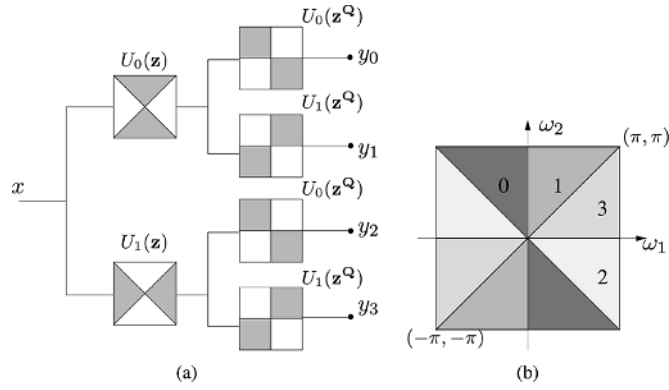
Fig. 3. Four-channel nonsubsampled directional filter bank constructed with two-channel fan filter banks. (a) Filtering structure. The equivalent filter in each channel is given by $U_k^{eq}(\mathbf{z}) = U_i(\mathbf{z})U_j(\mathbf{z}^{\mathbf{Q}})$. (b) Corresponding frequency decomposition.
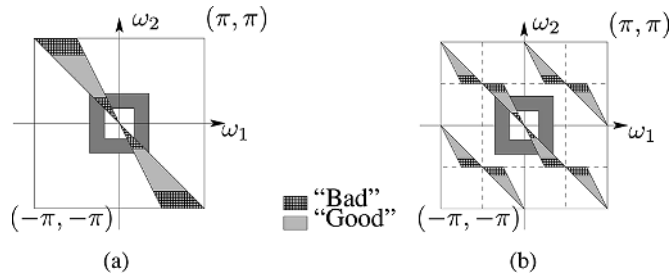

Fig. 4. Need for upsampling in the NSCT. (a) With no upsampling, the high-pass at higher scales will be filtered by the portion of the directional filter that has "bad" response. (b) Upsampling ensures that filtering is done in the "good" region.
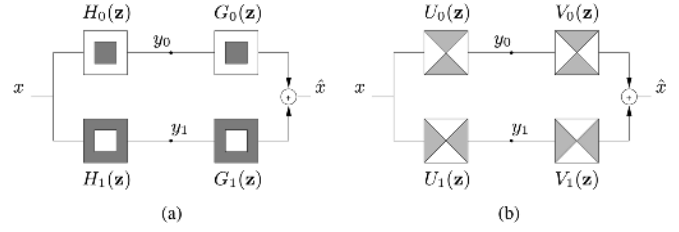

Fig. 5. Two-channel NSFBs used in the NSCT. The system is two times redundant and the reconstruction is error free when the filters satisfy Bezout's identity. (a) Pyramid NSFB. (b) Fan NSFB.

four channel decomposition. Note that in the second level, the upsampled fan filters $U_i(\mathbf{z}^{\mathbf{Q}})$, $i = 0, 1$ have checker-board frequency support, and when combined with the filters in the first level give the four directional frequency decomposition shown in Fig. 3. The synthesis filter bank is obtained similarly. Just like the critically sampled directional filter bank, all filter banks in the nonsubsampled directional filter bank tree structure are obtained from a single NSFB with fan filters [see Fig. 5(b)]. Moreover, each filter bank in the NSDFB tree has the same computational complexity as that of the building-block NSFB.

3) Combining the Nonsubsampled Pyramid and Nonsubsampled Directional Filter Bank in the NSCT: The NSCT is constructed by combining the NSP and the NSDFB as shown in Fig. 1(a). In constructing the NSCT, care must be taken when applying the directional filters to the coarser scales of the pyramid. Due to the tree-structure nature of the NSDFB, the directional response at the lower and upper frequencies suffers from aliasing which can be a problem in the upper stages of the pyramid (see Fig. 8). This is illustrated in Fig. 4(a), where the passband region of the directional filter is labeled as "Good" or "Bad." Thus, we see that for coarser scales, the high-pass channel in effect is filtered with the bad portion of the directional filter passband. This results in severe aliasing and in some observed cases a considerable loss of directional resolution.

We remedy this by judiciously upsampling the NSDFB filters. Denote the kth directional filter by $U_k(\mathbf{z})$. Then for higher scales, we substitute $U_k(\mathbf{z}^{2^m\mathbf{I}})$ for $U_k(\mathbf{z})$ where m is chosen

to ensure that the good part of the response overlaps with the pyramid passband. Fig. 4(b) illustrates a typical example. Note that this modification preserves perfect reconstruction. In a typical five-scale decomposition, we upsample by $2\mathbf{I}$ the NSDFB filters of the last two stages.

Filtering with the upsampled filters does not increase computational complexity. Specifically, for a given sampling matrix $\mathbf{S}$ and a 2-D filter $H(\mathbf{z})$, to obtain the output $y[\mathbf{n}]$ resulting from filtering $x[\mathbf{n}]$ with $H(\mathbf{z}^{\mathbf{S}})$, we use the convolution formula

$$y[\mathbf{n}] = \sum_{\mathbf{k} \in \text{supp}(h)} h[\mathbf{k}]x[\mathbf{n} - \mathbf{S}\mathbf{k}]. \tag{1}$$

This is the *à trous* filtering algorithm [9] ("à trous" is French for "with holes"). Therefore, each filter in the NSDFB tree has the same complexity as that of the building-block fan NSFB. Likewise, each filtering stage of the NSP has the same complexity as that incurred by the first stage. Thus, the complexity of the NSCT is dictated by the complexity of the building-block NSFBs. If each NSFB in both NSP and NSDFB requires L operations per output sample, then for an image of N pixels the NSCT requires about BNL operations where B denotes the number of subbands. For instance, if $L = 32$, a typical decomposition with four pyramid levels, 16 directions in the two finer scales, and eight directions in the two coarser scales would require a total of 1536 operations per image pixel.

If the building block 2-channel NSFBs in the NSP and NSDFB are invertible, then clearly the NSCT is invertible. It also underlies a frame expansion (see Section II-C). The frame elements are localized in space and oriented along a discrete set of directions. The NSCT is flexible in that it allows any number of 2^l directions in each scale. In particular, it can satisfy the anisotropic scaling law—a key property in establishing the expansion nonlinear approximation behavior [13], [14]. This property is ensured by doubling the number of directions in the NSDFB expansion at every other scale. The NSCT has redundancy given by $1 + \sum_{j=1}^{J} 2^{l_j}$, where l_j denotes the number of levels in the NSDFB at the jth scale.

B. Nonsubsampled Filter Banks

At the core of the proposed NSCT structure is the 2-D two-channel NSFB. Shown in Fig. 5 are the NSFBs needed to construct the NSCT. In this paper, we focus exclusively on the FIR case simply because it is easier to implement in multiple dimensions. For a general FIR two-channel NSFB, perfect reconstruc-

tion is achieved provided the filters satisfy the *Bezout identity*, as follows:

$$H_0(\mathbf{z})G_0(\mathbf{z}) + H_1(\mathbf{z})G_1(\mathbf{z}) = 1. \tag{2}$$

The Bezout relation in (2) puts no constraint on the frequency response of the filters involved. Therefore, to obtain good solutions, one has to impose additional conditions on the filters.

C. Frame Analysis of the NSCT

The NSFB can be interpreted in terms of analysis/synthesis operators of *frame systems*. A family of vectors $\{\phi_n\}_{n\in\Gamma}$ constitute a frame for a Hilbert space $\mathcal{H}$ if there exist two positive constants A, B such that for each $x \in \mathcal{H}$ we have

$$A\|x\|^2 \leq \sum_{n\in\Gamma} |\langle x, \phi_n\rangle|^2 \leq B\|x\|^2. \tag{3}$$

In the event that $A = B$, the frame is said to be tight. The *frame bounds* are the tightest positive constants satisfying (3).

Consider the NSFB of Fig. 5(a). The family $\{h_0[\cdot - \mathbf{n}], h_1[\cdot - \mathbf{n}]\}_{\mathbf{n}\in\mathbb{Z}^2}$ is a frame for $\ell_2(\mathbb{Z}^2)$ if and only if there exist constants $0 < A \leq B < \infty$ such that [21]

$$A \leq \underbrace{|H_0(e^{j\omega})|^2 + |H_1(e^{j\omega})|^2}_{t(e^{j\omega})} \leq B. \tag{4}$$

Thus, the frame bounds of an NSFB can be computed by

$$A = \text{ess.} \inf_{w\in[-\pi,\pi]^2} t(e^{j\omega}), \quad B = \text{ess.} \sup_{w\in[-\pi,\pi]^2} t(e^{j\omega}) \tag{5}$$

where ess. inf, and ess. sup denote the essential infimum and essential supremum respectively. From (4), we see that the frame is tight whenever $t(e^{j\omega})$ is almost everywhere constant. For FIR filters, this means that $H_0(\mathbf{z})H_0(\mathbf{z}^{-1}) + H_1(\mathbf{z})H_1(\mathbf{z}^{-1}) = c$. Such condition can only be met with linear phase FIR filters if $H_0(\mathbf{z})$ and $H_1(\mathbf{z})$ are either trivial delays or combination of two delays (for a formal proof, see [22, pp. 337–338] or [23]).

Because the NSFB is redundant, an infinite number of inverses exist. Among them, the *pseudoinverse* is optimal in the least-square sense [1]. Given a frame of analysis filters, the synthesis filters corresponding to the frame pseudoinverse are given by $G_i(\mathbf{z}) = \hat{H}_i(\mathbf{z})/t(\mathbf{z})$ for $i = 0, 1$ [21]. In this case, the synthesis filters form the *dual frame* with lower and upper frame bounds given by B^{-1} and A^{-1} respectively. When the analysis filters are FIR, then unless the frame is tight, the synthesis filters of the pseudoinverse will be IIR.

From the above discussion, we gather two important points: 1) linear phase filters and tight frames are mutually exclusive and 2) the pseudoinverse is desirable, but is IIR if the frame is not tight. Consequently, an FIR NSFB system with linear phase filters and with synthesis filters corresponding to the pseudoinverse is not possible. However, we can approximate the pseudoinverse with FIR filters. For a given number of filter taps, the closer the frame is to being tight, the better an FIR approximation of the pseudoinverse can be [24]. Thus, in the designs of

J	A actual	A estimated	B actual	B estimated
1	0.9586	0.9596	1.0435	1.0435
2	0.9393	0.9189	1.0504	1.0889
3	0.9332	0.8808	1.0515	1.1362
4	0.9316	0.8444	1.0517	1.1857

the filters we seek linear phase filters that underly a frame that is as close to a tight one as possible.

In a general FIR perfect reconstruction NSFB system, both analysis and synthesis filters form a frame. If we denote the analysis and synthesis frame bounds by A_a, B_a and A_s, B_s respectively, the frames will be close to tight provided [24]

$$r_a := A_a/B_a \approx 1, \text{ and } r_s := A_s/B_s \approx 1.$$

We always assume that the filters are normalized so that we have $A_a \leq 1 \leq B_a$. In case the pseudoinverse is used, then we also have $A_s \leq 1 \leq B_s$. The following result shows the NSCT is a frame operator for $\ell_2(\mathbf{Z}^2)$ whenever the constituent NSFBs each forms a frame.

Proposition 1: In the NSCT, if the pyramid filter bank constitute a frame with frame bounds A_p and B_p, and the fan filters constitute a frame with frame bounds A_q and B_q, then the NSCT is a frame with bounds A and B, satisfying

$$A_p^J A_q^{\min\{l_j\}} \leq A \leq B \leq B_p^J B_q^{\max\{l_j\}}.$$

Proof: See *Appendix A:*

Remark 1: When both the pyramid and fan filter banks form tight frames with bound 1, then $A_p = A_q = B_p = B_q = 1$, and from the above proposition, the NSCT is also a tight frame with bound 1.

The above estimates on A and B can be accurate in some cases, especially when the frame is close to a tight one and the number of levels is small (e.g., $J = 4$). In general, however, they are not accurate estimates. Their purpose is more of giving an interval for the frame bounds rather than the actual values. Table I shows estimates for different number of scales. The actual frame bound is computed from (4)–(5), whereas the estimates are the ones given according to Proposition 1.

III. FILTER DESIGN AND IMPLEMENTATION

The filter-design problem of the NSCT comprises the two basic NSFBs displayed in Fig. 5. The goal is to design the filters imposing the Bezout identity (i.e., perfect reconstruction) and enforcing other properties such as sharp frequency response, easy implementation, regularity of the frame elements, and tightness of the corresponding frames. It is also desirable that the filters are linear-phase.

Two-channel 1-D NSFBs that underly tight frames are designed in [21]. However, the design methodology of [21] is not easy to extend to 2-D designs since it relies on spectral factorization which is hard in 2-D. If we relax the tightness constraint,
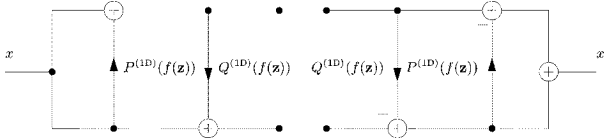
Fig. 6. Lifting structure for the NSFB designed with the mapping approach. The 1-D prototype is factored with the Euclidean algorithm. The 2-D filters are obtained by replacing $x \mapsto f(\mathbf{z})$.

then the design becomes more flexible. In addition, as we alluded to earlier, non-tight filters can be linear phase.

An effective and simple way to design 2-D filters is the *mapping* approach first proposed by McClellan [25] in the context of digital filters and then used by several authors [23], [26]–[28] in the context of filter banks. In such an approach, the 2-D filters are obtained from 1-D ones. In the context of NSFBs, a set of perfect reconstruction 2-D filters is obtained in the following way.

Step 1) Construct a set of 1-D polynomials $\{H_i^{(1D)}(x), G_i^{(1D)}(x)\}_{i=0,1}$ that satisfies the Bezout identity.

Step 2) Given a 2-D FIR filter $f(\mathbf{z})$, then $\{H_i^{(1D)}(f(\mathbf{z})), G_i^{(1D)}(f(\mathbf{z}))\}_{i=0,1}$ are 2-D filters satisfying the Bezout identity.

Thus, one has to design the set of 1-D filters and the mapping function $f(\mathbf{z})$ so that the ideal responses are well approximated with a small number of filter coefficients. In the mapping approach, one can control the frequency and phase responses through the mapping function. Moreover, if the mapping function is zero-phase, then $f(\mathbf{z}) = f(\mathbf{z}^{-1})$ and it follows that the mapped filters are also zero-phase. In this case, on the unit sphere, the mapping function is a 2-D polynomial in $(\cos \omega_1, \cos \omega_2)$. We thus denote it by $F(x_1, x_2)$, where it is implicit that $f(e^{j\omega}) = F(\cos \omega_1, \cos \omega_2)$.

A. Implementation Through Lifting

Filters designed with the mapping approach can be efficiently factored into a ladder [29] or lifting [30] structure that simplifies computations. To see this, assume without loss of generality that the degree of the high-pass prototype polynomial $H_1^{(1D)}(x)$ is smaller than that of $H_0^{(1D)}(x)$. Suppose also that there are synthesis filters $G_0^{(1D)}(x)$ and $G_1^{(1D)}(x)$ such that the Bezout identity is satisfied. In this case it follows that $\gcd\{H_0^{(1D)}, H_1^{(1D)}\} = 1$. The Euclidean algorithm then enables us to factor the filters in the following way [29], [30]:

$$\begin{pmatrix} H_0^{(1D)}(x) \\ H_1^{(1D)}(x) \end{pmatrix} = \prod_{i=0}^{N} \begin{pmatrix} 1 & 0 \\ P_i^{(1D)}(x) & 1 \end{pmatrix}$$
$$\times \begin{pmatrix} 1 & Q_i^{(1D)}(x) \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix}. \quad (6)$$

As a result, we can obtain a 2-D factorization by replacing x with $f(\mathbf{z})$. This factorization characterizes *every* 2-D NSFB derived from 1-D through the mapping method. Fig. 6 illustrates the ladder structure with one stage.

In general, the lifting implementation at least halves the number of multiplications and additions of the direct form [30].

This can be shown by following the same reasoning as in the critically-sampled case discussed in [30]. The complexity can be reduced further if the lifting steps in the 1-D prototype are monomials and the mapping filter $f(\mathbf{z})$ has the form

$$f(\mathbf{z}) = f_1(z_1^{p_1} z_2^{p_2}) f_2(z_1^{q_1} z_2^{q_2}) \quad (7)$$

for suitable $f_1(z)$, $f_2(z)$, and integers p_1, p_2, q_1, q_2. Note that if $f(z)$ is a 1-D filter, then the 2-D filter $f(z_1^p z_2^q)$ for integers p and q has the same complexity as that of $f(z)$. Therefore, filters of the form in (7) have the same complexity as that of separable filters (i.e., filters of the form $f(\mathbf{z}) = f(z_1)f(z_2)$) which amounts to two 1-D filtering operations. Notice that if $f(\mathbf{z})$ is as in (7), then for an arbitrary sampling matrix $\mathbf{S}$, $f(\mathbf{z}^{\mathbf{S}})$ also has the form in (17). Consequently, all NSFBs in the NSDFB tree structure can be implemented with 1-D operations whenever the building block fan NSFB can be implemented with 1-D filtering operations. The same reasoning applies to the NSFBs of the NSP.

B. Pyramid Filter Design

In the pyramid case, we impose *line zeros* at $\omega = (\pi, \omega_2)$ and $\omega = (\omega_1, \pi)$. Notice that an Nth order line zero at $\omega = (\pi, \omega_2)$, for example, amounts to a $(1 + e^{j\omega_1})^N$ factor in the low-pass filter. Such zeros are useful to obtain good approximation of the ideal frequency response of Fig. 5, in addition to imposing regularity of the scaling function. We point out that for the approximation of smooth C^α images, $1 + \lfloor \alpha \rfloor$ point zeros at $(\pm \pi, \pm \pi)$ would suffice [31]. However, our experience shows that point zeros alone do not guarantee a "reasonable" frequency response of the pyramid filters. The following proposition characterizes the mapping function that generates zeros in the resulting 2-D filters.

Proposition 2: Let $G^{(1D)}(z)$ be a polynomial with roots $\{z_i\}_{i=1}^n$ where each z_i has multiplicity n_i. Suppose we want a mapping function $F(x_1, x_2)$ such that

$$G^{(1D)}(F(x_1, x_2)) = (x_1 - c)^{N_1} (x_2 - d)^{N_2} L(x_1, x_2) \quad (8)$$

where $L(x_1, x_2)$ is a bivariate polynomial. Then $G^{(1D)}(F(x_1, x_2))$ has the form in (8) *if and only if* $F(x_1, x_2)$ takes the form

$$F(x_1, x_2) = z_j + (x_1 - c)^{N_1'} (x_2 - d)^{N_2'} L_F(x_1, x_2) \quad (9)$$

for some root $z_j \in \{z_i\}_{i=1}^n$, where $L_F(x_1, x_2)$ is a bivariate polynomial, and N_1', N_2' are such that $N_1' n_i \geq N_1$ and $N_2' n_j \geq N_2$.

Proof: See *Appendix B:*

The above result holds in general, even for point zeros (as opposed to line zeros as in Proposition 2). We will explore this extensively in the designs that follow.

Suppose the prototype filters $H_0^{(1D)}(x), G_0^{(1D)}(x)$ each has zeros at $x = 1$. Then, in order to produce a suitable zero-phase mapping function for the pyramid NSFB we consider the class of maximally-flat filters given by the polynomials [32, p. 171]

$$P_{N,L}(x) := \left(\frac{1+x}{2}\right)^N \sum_{l=0}^{L-1-N} \binom{N+l-1}{l} \left(\frac{1-x}{2}\right)^l \quad (10)$$

TABLE II
MAXIMALLY FLAT MAPPING POLYNOMIALS USED IN THE
DESIGN OF THE NONSUBSAMPLED FAN FILTER BANK

N	$F_N^{(\text{fan})}(x_1, x_2)$
1	$\frac{1}{2}(x_1 - x_2)$
2	$\frac{1}{4}(x_1 - x_2)(3 + x_1 x_2)$
3	$\frac{1}{16}(x_1 - x_2)(15 - x_1^2 + 8x_1 x_2 - x_2^2 + 3x_1^2 x_2^2)$
4	$\frac{1}{32}(x_1 - x_2)(35 - 5x_1^2 + 25x_1 x_2 - 3x_1^3 x_2 - 5x_2^2 + 15x_1^2 x_2^2 - 3x_1 x_2^3 + 5x_1^3 x_2^3)$
5	$\frac{1}{256}(x_1 - x_2)\,(315 - 70x_1^2 + 3x_1^4 + 280x_1 x_2 - 72x_1^3 x_2 - 70x_2^2 + 228x_1^2 x_2^2$ $\quad - 30x_1^4 x_2^2 - 72x_1 x_2^3 + 120x_1^3 x_2^3 + 3x_2^4 - 30x_1^2 x_2^4 + 35x_1^4 x_2^4)$

where N controls the degree of flatness at $x = -1$ and L controls the degree flatness at $x = 1$. Following Proposition 2, we can construct a family of mapping functions as

$$F^{(\text{pyr})}(x_1, x_2) = -1 + 2P_{N_1, L_1}(x_1)P_{N_2, L_2}(x_2) \quad (11)$$

so that zero moments at $x_1 = -1$ and $x_2 = -1$ are guaranteed. Note that, except for the constant -1, $F(x_1, x_2)$ has the form of (7) and hence can be implemented with 1-D filtering operations only.

C. Fan Filter Design

To design the fan filters idealized in Fig. 5(b), we use the same methodology as in the pyramid case. The distinction occurs in the mapping function. The fan-shaped response can be obtained from a diamond-shaped response by simple modulation in one of the frequency variables. This modulation preserves the perfect reconstruction property. A useful family of mapping functions for the diamond-shaped response is obtained by imposing flatness around $\omega = (\pm\pi, \pm\pi)$ and $\omega = (0, 0)$ in addition to point zeros at $(\pm\pi, \pm\pi)$. If the mapping function is zero-phase, this amounts to imposing flatness in a polynomial at $(x_1, x_2) = (-1, -1)$ and $(x_1, x_2) = (1, 1)$, and zeros at $(x_1, x_2) = (-1, -1)$. To obtain the fan mapping function, we simply replace x_2 with $-x_2$, which corresponds to a shift by π on the ω_2 variable. If we use prototype 1-D filters with zeros at $x = 1$, then similar to the pyramid case, we set

$$F_N^{(\text{fan})}(x_1, x_2) = -1 + Q_N(x_1, -x_2) \quad (12)$$

where the polynomials $Q_N(x_1, x_2)$ give the class of maximally-flat half-band filters with diamond support. A closed-form expression for $Q_N(x_1, x_2)$ is given in [33]. Table II displays the first six mapping functions $F_N^{(\text{fan})}(x_1, x_2)$ for the resulting maximally-flat fan filter bank.

D. Design Examples

The design through mapping is based on a set of 1-D polynomials that satisfies the Bezout identity, i.e., $H_0^{(1D)}(x)G_0^{(1D)}(x) + H_1^{(1D)}(x)G_1^{(1D)}(x) = 1$. The design can be simplified if we impose the restriction that $H_1^{(1D)}(x) = G_0^{(1D)}(-x)$ and $G_1^{(1D)}(x) = H_0^{(1D)}(-x)$. One advantage of this choice is that the frequency response of the filters can be controlled by the low-pass branch—the high-pass will automatically have the complementary response. Another advantage is that, under an

additional condition, the frame bounds of the analysis and synthesis frames are the same and can be computed from the 1-D prototypes. To see this, suppose that f is the mapping function and that $\mathbf{Ran}(f) = \mathbf{Ran}(-f)$ with $\mathbf{Ran}(f)$ denoting the range of the mapping function f. Then using (5), we have

$$A_a = \inf_{\omega \in [\pi, \pi]^2} |H_0^{(1D)}(f(e^{j\omega}))|^2 + |G_0^{(1D)}(-f(e^{j\omega}))|^2$$

$$= \inf_{x \in \mathbf{Ran}(f)} |H_0^{(1D)}(x)|^2 + |G_0^{(1D)}(-x)|^2 \quad (13)$$

$$= \inf_{x \in \mathbf{Ran}(-f)} |H_0^{(1D)}(-x)|^2 + |G_0^{(1D)}(x)|^2$$

$$= \inf_{x \in \mathbf{Ran}(f)} |H_0^{(1D)}(-x)|^2 + |G_0^{(1D)}(x)|^2 = A_s. \quad (14)$$

A similar argument shows that

$$B_a = B_s = \sup_{x \in \mathbf{Ran}(f)} |H_0^{(1D)}(-x)|^2 + |G_0^{(1D)}(x)|^2. \quad (15)$$

Example 1 (Pyramid Filters Very Close to Tight Ones): In order to get filters that are almost tight, we design the prototypes $H_0^{(1D)}(x)$ and $G_0^{(1D)}(x)$ to be very close to each other. If we let $H_1^{(1D)}(x) = G_0^{(1D)}(-x)$ and $G_1^{(1D)}(x) = H_0^{(1D)}(-x)$, then the following filters can be checked to satisfy the Bezout identity:

$$H_0^{(1D)}(x) = K(1 + k_2 x + k_2 k_3 x^2)$$
$$G_0^{(1D)}(x) = K^{-1}(1 - k_1 x - k_3 x + k_1 k_2 x^2 - k_1 k_2 k_3 x^3).$$

The above filters are obtained from a ladder structure where the ladder steps are given by $k_j x$, $j = 1, 2, 3$. To obtain the prototype, we choose the constants $K, k_1, k_2,$ and k_3 such that each filter has a zero at $x = 1$ and, in addition, that $H_0^{(1D)}(-1) = G_0^{(1D)}(1) = 1$. We obtain

$$H_0^{(1D)}(x) = \frac{1}{2}(x + 1)\left(\sqrt{2} + (1 - \sqrt{2})x\right),$$

$$G_0^{(1D)}(x) = \frac{1}{2}(x + 1)\left(\sqrt{2} + (4 - 3\sqrt{2})x + (2\sqrt{2} - 3)x^2\right).$$

The lifting factorization of the prototype filters is given by

$$\begin{pmatrix} H_0^{(1D)}(x) \\ H_1^{(1D)}(x) \end{pmatrix} = \begin{pmatrix} 1 & \alpha x \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \beta x & 1 \end{pmatrix} \begin{pmatrix} 1 & \gamma x \\ 0 & 1 \end{pmatrix} \begin{pmatrix} K_1 \\ K_2 \end{pmatrix} \quad (16)$$

with

$$\alpha = \gamma = 1 - \sqrt{2}, \quad \beta = \frac{1}{\sqrt{2}}, \quad K_1 = K_2 = \frac{1}{\sqrt{2}}.$$

Notice that this implementation has four multiplies/sample, whereas the direct form has seven multiplies/sample.

In this example we set $F^{(\text{pyr})}(x_1, x_2) = -1 + 2P_{2,4}(x_1)P_{2,4}(x_2)$, where $P_{N,L}$ is given in (11) so that each of the filters has a fourth-order zero at $\omega_1 = \pm\pi$ and at $\omega_2 = \pm\pi$. Since the ladder steps are monomials, the NSFB can be implemented with 1-D filtering operations. The frame bounds are computed using (13)–(15)

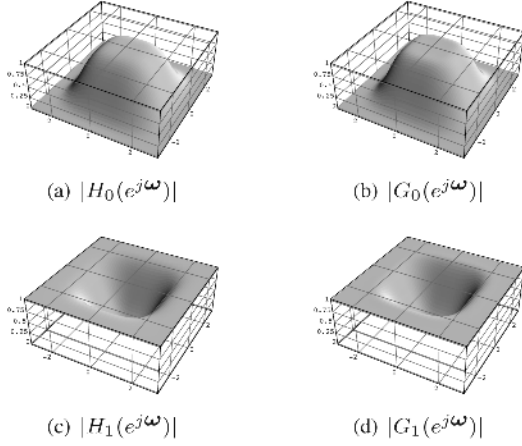$$A_a = A_s = 0.96, \quad B_a = B_s = 1.04.$$

(a) $|H_0(e^{j\boldsymbol{\omega}})|$ (b) $|G_0(e^{j\boldsymbol{\omega}})|$

(c) $|H_1(e^{j\boldsymbol{\omega}})|$ (d) $|G_1(e^{j\boldsymbol{\omega}})|$

Fig. 7. Magnitude response of the filters designed in Example 1 with maximally-flat filters. The nonsubsampled pyramid filter bank underlies almost tight analysis and synthesis frames.

(a) $|U_0(e^{j\boldsymbol{\omega}})|$ (b) $|V_0(e^{j\boldsymbol{\omega}})|$

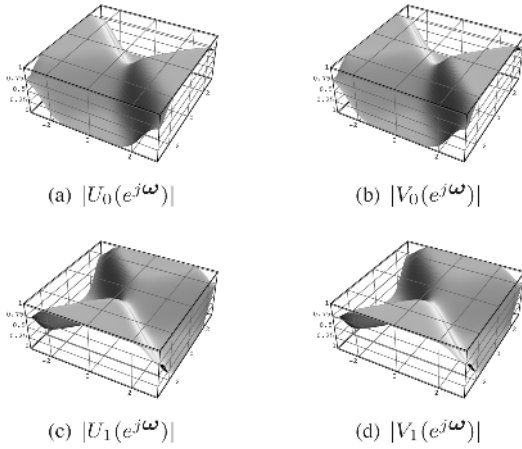(c) $|U_1(e^{j\boldsymbol{\omega}})|$ (d) $|V_1(e^{j\boldsymbol{\omega}})|$

Fig. 8. Fan filters designed with prototype filters of Example 1 and the fan mapping function of Table II.

Thus, we have $r_a = r_s = 1.083$ and the frame is almost tight. The support size of $H_0(\mathbf{z})$ is 13×13 whereas $G_0(\mathbf{z})$ has support size 19×19. Fig. 7 shows the response of the resulting filters.

Example 2 (Maximally-Flat Fan Filters): Using the prototype 1-D filters of Example 1, we choose $F_N^{(\mathrm{fan})}(x_1, x_2)$ in Table II with $N = 3$. We then replace x in (16) with $F_N^{(\mathrm{fan})}(x_1, x_2)$ to obtain $U_0(\mathbf{z})$ and $V_0(\mathbf{z})$. Fig. 8 shows the magnitude response of the filters. The support size of $U_0(\mathbf{z})$ is 21×21, whereas the support size of $V_0(\mathbf{z})$ is 31×31.

E. Regularity of the NSCT Basis Functions

The regularity of the NSCT basis functions can be controlled by the NSP low-pass filters. Denoting by $H_0(\omega)$ the scaling low-pass filter used in the pyramid (we write $H_0(\omega)$ instead of $H_0(e^{j\omega})$ for convenience), we have the associated scaling function

$$\Phi(\omega :) = \prod_{j=1}^{\infty} H_0(2^{-j}\omega)$$

where convergence is in the weak sense. In our proposed design the filter $H_0(\omega)$ can be factored as

$$H_0(\omega) = \left(\frac{1 + e^{jw_1}}{2}\right)^{N_1} \left(\frac{1 + e^{jw_2}}{2}\right)^{N_2} R_{H_0}(\omega). \quad (17)$$

Notice that the remainder filter $R_{H_0}(\omega)$ is not separable. Therefore, one cannot separate the regularity estimation as two 1-D problems. Nonetheless, a similar argument can be developed and an estimate of the 2-D regularity of the scaling filter is obtained.

Proposition 3: Let $H_0(\omega)$ be a scaling filter as in (17) with the corresponding scaling function $\Phi(\omega)$. Let

$$B = \sup_{\omega \in [\pi, \pi]^2} |R_{H_0}(\omega)|.$$

Then

$$|\Phi(\omega)| \leq C \left(\frac{1}{1 + |w_1|}\right)^{N_1 - \log_2 B} \left(\frac{1}{1 + |w_2|}\right)^{N_2 - \log_2 B}.$$

Proof: See *Appendix C:*

As an example, consider the prototype filter in Example 1 and the mapping $F^{(\mathrm{pyr})}(x_1, x_2) = -1 + 2P_{1,2}(x_1)P_{1,2}(x_2)$. The resulting filter has second-order zeros at $\omega_1 = \pm\pi$ and at $\omega_2 = \pm\pi$. It can be verified that $|R_{H_0}(\omega)| \lessapprox 1.83$ and $|R_{G_0}(\omega)| \lessapprox 1.49$ so that the regularity exponent[2] is at least $2 - \log_2 1.83 \approx 1.13$ for $H_0(\omega)$ and $2 - \log_2 1.49 \approx 1.43$ for $G_0(\omega)$. Thus the corresponding scaling functions and wavelets are at least continuous. We point out that better estimates are possible applying similar 1-D techniques. For instance, one could prove a result similar to Lemma 7.1.2 in [32, p. 217], as a consequence of Proposition 3.

Fig. 9 shows the basis functions of the NSCT obtained with the filters designed via mapping. As the picture shows, the functions have a good degree of regularity.

IV. APPLICATIONS

A. Image Denoising

In order to illustrate the potential of the NSCT designed using the techniques previously discussed, we study additive white Gaussian noise (AWGN) removal from images by means of thresholding estimators.

1) Comparison to Other Transforms: To highlight the performance of the NSCT relative to other transforms, we perform hard threshold on the subband coefficients of the various transforms. We choose the threshold

$$T_{i,j} = K\sigma_{n_{ij}}$$

for each subband. This has been termed K-sigma thresholding in [34]. We set $K = 4$ for the finest scale and $K = 3$ for the remaining ones. We use five scales of decomposition for both

---

[2]The regularity exponent of a scaling function $\phi(\mathbf{t})$ is the largest number α such that $\hat{\Phi}(\omega)$ decays as fast as $1/(1 + |\omega_1| + |\omega_2|)^\alpha$.
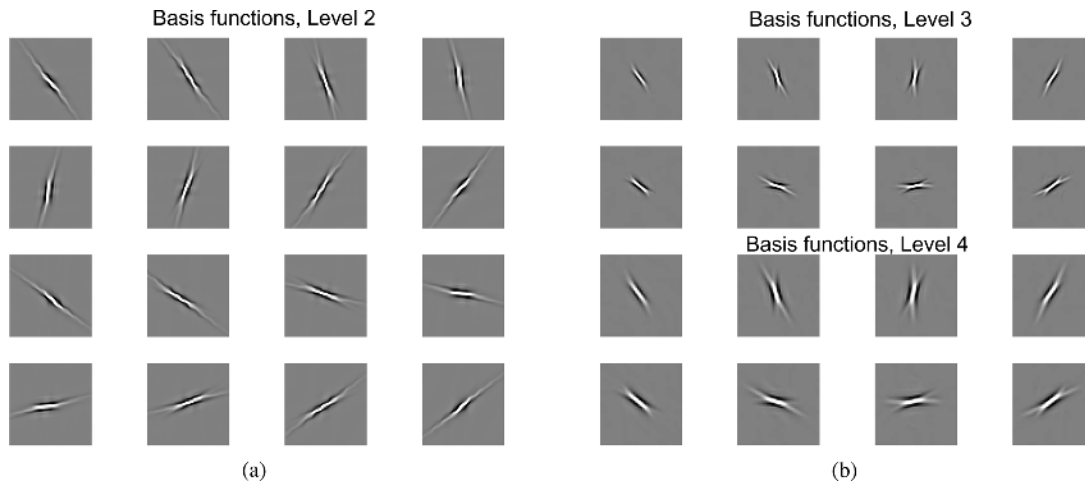
Fig. 9. Basis functions of the nonsubsampled contourlet transform. (a) Basis functions of the second stage of the pyramid. (b) Basis functions of third (top eight) and fourth (bottom eight) stages of the pyramid.

TABLE III
DENOISING PERFORMANCE OF THE NSCT. THE LEFT-MOST COLUMNS ARE HARD THRESHOLDING AND THE RIGHT-MOST ONES SOFT ESTIMATORS. FOR HARD THRESHOLDING, THE NSCT CONSISTENTLY OUTPERFORMS CURVELETS AND THE NSWT. THE NSCT-LAS PERFORMS ON A PAR WITH THE MORE SOPHISTICATED ESTIMATOR BLS-GSM [8] AND IS SUPERIOR TO THE BIVSHRINK ESTIMATOR OF [7]

Lena	Comparison to other transforms					Comparison to other methods			
σ	Noisy	NSWT	Curvelets	CT	NSCT	NSWT-LAS	BivShrink	BLS-GSM	NSCT-LAS
10	28.13	34.26	34.17	31.90	**34.69**	35.19	35.34	**35.59**	35.46
20	22.13	31.40	31.52	28.34	**32.03**	32.12	32.40	**32.62**	32.50
30	18.63	29.66	30.01	27.10	**30.35**	30.30	30.54	**30.84**	30.70
40	16.13	28.37	28.84	25.84	**29.10**	29.01	-	**29.58**	29.38
50	14.20	27.41	27.78	24.87	**28.10**	28.00	-	**28.61**	28.34

Barbara	Comparison to other transforms					Comparison to other methods			
σ	Noisy	NSWT	Curvelets	CT	NSCT	NSWT-LAS	BivShrink	BLS-GSM	NSCT-LAS
10	28.17	31.58	32.28	29.62	**33.01**	33.40	33.35	34.03	**34.09**
20	22.15	27.23	28.89	26.26	**29.41**	29.45	29.80	30.28	**30.60**
30	18.63	25.10	26.93	24.42	**27.24**	27.22	27.65	28.11	**28.56**
40	16.14	24.02	25.51	23.16	**25.79**	25.76	-	26.58	**27.12**
50	14.20	23.37	24.31	22.29	**24.79**	24.72	-	25.43	**26.02**

Peppers	Comparison to other transforms					Comparison to other methods		
σ	Noisy	NSWT	Curvelets	CT	NSCT	NSWT-LAS	BLS-GSM	NSCT-LAS
10	28.17	33.71	33.59	31.30	**33.81**	34.35	**34.63**	34.41
20	22.15	31.19	31.13	28.57	**31.60**	31.65	**32.06**	31.82
30	18.63	29.43	29.45	26.81	**30.07**	29.95	**30.41**	30.19
40	16.13	28.09	28.01	25.50	**28.85**	28.65	**29.20**	28.95
50	14.20	27.04	26.70	24.47	**27.82**	27.62	**28.24**	27.93

NSCT, contourlet transform (CT), and NSWT. For the NSCT and CT we use 4, 8, 8, 16, 16 directions in the scales from coarser to finer, respectively.

Table III (left columns) shows the PSNR results for various transforms and noise intensities. The results show the NSCT is *consistently superior* to curvelets and NSWT in PSNR measure. For the "Barbara" image the NSCT yields improvements in ex-

cess of 1.90 dB in PSNR over the NSWT. The NSCT also is superior to the CT as the results show. Fig. 10 displays the reconstructed images using the the NSWT, curvelets and NSCT. As the figure shows, both the NSCT and the curvelet transform offers a better recovery of edge information relative to the NSWT. But improvements can be seen in the NSCT, particularly around the eye.

Fig. 10. Image denoising with the NSCT and hard thresholding. The noisy intensity is 20. (a) Original Lena image. (b) Denoised with the NSWT, $\mathrm{PSNR} = 31.40$ dB. (c) Denoised with the curvelet transform and hard thresholding, $\mathrm{PSNR} = 31.52$ dB. (d) Denoised with the NSCT, $\mathrm{PSNR} = 32.03$ dB.

2) Comparison To Other Denoising Methods: We perform soft thresholding (shrinkage) independently in each subband. Following [6], we choose the threshold

$$T_{i,j} = \sigma_{N_{ij}}^2 / \sigma_{i,j,\mathbf{n}}$$

where $\sigma_{i,j,\mathbf{n}}$ denotes the variance of the $\mathbf{n}$-th coefficient at the i-th directional subband of the j-th scale, and $\sigma_{N_{ij}}^2$ is the noise variance at scale j and direction i. It is shown in [6] that shrinkage estimation with $T = (\sigma^2/\sigma_X)$, and assuming X generalized Gaussian distributed yields a risk within 5% of the optimal Bayes risk. As studied in [35], contourlet coefficients are well modelled by generalized Gaussian distributions. The signal variances are estimated locally using the neighboring coefficients contained in a square window within each subband and a maximum likelihood estimator. The noise variance in each subband is inferred using a Monte Carlo technique where the variances are computed for a few normalized noise images and then averaged to stabilize the results. We refer to this method as *local adaptive shrinkage* (LAS). Effectively, our LAS method is a simplified version of the denoising method proposed in [36] that works in the NSCT or NSWT domain. In the LAS estimator we use four scales for both the NSCT and NSWT. For the NSCT we use 3, 3, 4, 4 directions in the scales from coarser to finer, respectively.

To benchmark the performance of the NSCT-LAS scheme, we have used two of the best denoising methods in the literature: 1) bivariate shrinkage with local variance estimation (BivShrink) [7]; 2) Bayes least squares with a Gaussian scale-mixture model (BLS-GSM) proposed in [8]. Table III

(right columns) shows the results obtained.[3] The NSCT coupled with the LAS estimator (NSCT-LAS) produced very satisfactory results. In particular, among the methods studied, the NSCT-LAS yields the best results for the "Barbara" image, being surpassed by the BLS-GSM method for the other images. Despite its slight loss in performance relative to BLS-GSM, we believe the NSCT has potential for better results. This is because by comparison, the BLS-GSM is a considerably richer and more sophisticated estimation method than our simple local thresholding estimator. However, studying more complex denoising methods in the NSCT domain is beyond the scope of the present paper. Fig. 11 displays the denoised images with both BLS-GLM and NSCT-LAS methods. As the pictures show, the NSCT offers a slightly better reconstruction. In particular, the tablecloth texture is better recovered in the NSCT-LAS scheme.

We briefly mention that in denoising applications, one can reduce the redundancy of the NSCT by using critically sampled directional filter banks over the nonsubsampled pyramid. This results in a transform with $J + 1$ redundancy which is considerably faster. There is however a loss in performance as Table IV shows. Nonetheless, in some applications, the small performance loss might be a good price to pay given the reduced redundancy of this alternative construction.

B. Image Enhancement

Existing image-enhancement methods amplify noise when they amplify weak edges since they cannot distinguish noise from weak edges. In the frequency domain, both weak edges and noise produce low-magnitude coefficients. Since weak edges are geometric structures and noise is not, we can use the NSCT to distinguish them.

The NSCT is shift-invariant so that each pixel of the transform subbands corresponds to that of the original image in the same spatial location. Therefore, we gather the geometrical information pixel by pixel from the NSCT coefficients. We observe that there are three classes of pixels: strong edges, weak edges, and noise. First, the strong edges correspond to those pixels with large magnitude coefficients in all subbands. Second, the weak edges correspond to those pixels with large magnitude coefficients in some directional subbands but small magnitude coefficients in other directional subbands within the same scale. Finally, the noise corresponds to those pixels with small magnitude coefficients in all subbands. Based on this observation, we can classify pixels into three categories by analyzing the distribution of their coefficients in different subbands. One simple way is to compute the mean (denoted by mean) and the maximum (denoted by max) magnitude of the coefficients for each pixel across directional subbands, and then classify it by

$$\begin{cases} \text{strong edge,} & \text{if} \quad \text{mean} \geq c\sigma \\ \text{weak edge,} & \text{if} \quad \text{mean} < c\sigma, \max \geq c\sigma \\ \text{noise,} & \text{if} \quad \text{mean} < c\sigma, \max < c\sigma \end{cases} \quad (18)$$

where c is a parameter ranging from 1 to 5, and σ is the noise standard deviation of the subbands at a specific pyramidal level.

---

[3]The PSNR values of the BivShrink method were obtained from the tables in [7]. In [7], the authors do not use the "Peppers" image as a test image, hence we do not have a BivShrink column for "Peppers."
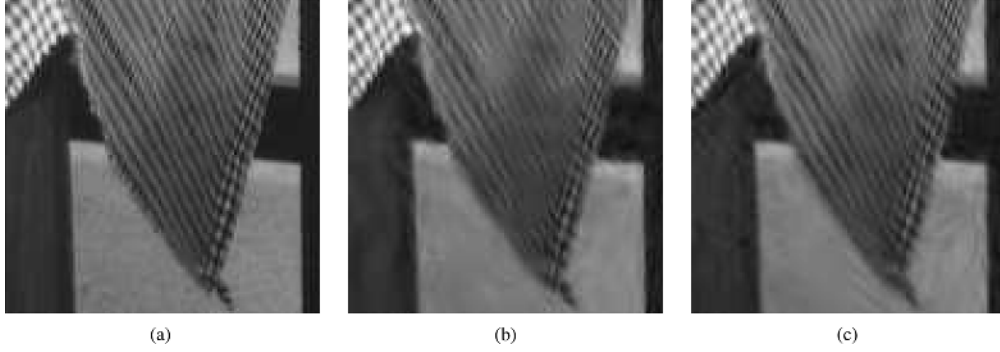
Fig. 11. Comparison between the NSCT-LAS and BLS-GSM denoising methods. The noise intensity is 20. (a) Original Barbara image. (b) Denoised with the BLS-GSM method of [8], $\mathrm{PSNR} = 30.28$ dB. (c) Denoised with NSCT-LAS, $\mathrm{PSNR} = 30.60$ dB.

TABLE IV
RELATIVE LOSS IN PSNR PERFORMANCE (dB) WHEN USING THE CRITICALLY SAMPLED DFB INSTEAD OF THE NSDFB IN THE NSCT-LAS METHOD

σ	"Lena"	"Barbara"	"Peppers"
10	-0.23	-0.48	-0.32
20	-0.09	-0.44	-0.09
30	-0.05	-0.37	-0.01
40	-0.04	-0.34	-0.00
50	-0.05	-0.31	-0.01

TABLE V
DV AND BV COMPARISONS FOR ENHANCEMENT METHODS

Image	Original		NSWT		NSCT	
	DV	BV	DV	BV	DV	BV
Barbara	12.56	0.63	14.72	0.74	35.39	0.52
Lena	6.40	0.85	7.65	0.72	16.45	0.63
Peppers	5.14	0.86	14.52	0.67	23.70	0.68
Boat	9.45	0.86	11.75	0.68	22.11	0.59

We first estimate the noise variance of the input image with the robust median operator [36] and then compute the noise variance of each subband [34].

The goal of image enhancement is to amplify weak edges and to suppress noise. To this end, we modify the NSCT coefficients according to the category of each pixel by a nonlinear mapping function (similar to [37])

$$y(x) = \begin{cases} x, & \text{strong edge pixels} \\ \max((\frac{c\sigma}{|x|})^p, 1)x, & \text{weak edge pixels} \\ 0, & \text{noise pixels} \end{cases} \quad (19)$$

where the input x is the original coefficient, and $0 < p < 1$ is the amplifying gain. This function keeps the coefficients of strong edges, amplifies the coefficients of weak edges, and zeros the noise coefficients.

We summarize our enhancement method using the NSCT in the following algorithm.

1) Compute the NSCT of the input image for N levels.
2) Estimate the noise standard deviation of the input image.
3) For each level of the pyramid:
   a) Estimate the noise variance.
   b) At each pixel location, compute the mean and the maximum magnitude of the corresponding coefficients in all directional subbands at this level, and classify each pixel according to (18) into "strong edges," "weak edges," or noise.
   c) For each directional subband, use the nonlinear mapping function given in (19) to modify the NSCT coefficients according to the classification.
4) Reconstruct the enhanced image from the modified NSCT coefficients.

We compare the enhancement results by the proposed algorithm with those by the NSWT. In the experiments, we choose $c = 4$ and $p = 0.3$. To evaluate the enhancement performance objectively, the *detailed variance* (DV) and *background variance* (BV) were proposed in [38]. The DV and BV values represent the variance of foreground and background pixels, respectively. A good enhancement methods should increase the DV of the original image but not the BV. We use the BV and DV to compare the enhancement performance of the NSWT and the NSCT in Table V. Fig. 12 shows the results obtained for the "Barbara" image. We observe that our proposed algorithm offers better results in enhancing the weak edges in the textures.

V. CONCLUSION

We have developed a fully shift-invariant version of the contourlet transform, the NSCT. The design of the NSCT is reduced to the design of a nonsubsampled pyramid filter bank and a nonsubsampled fan filter bank. We exploit this new less stringent filter-design problem using a mapping approach, thus dispensing with the need for 2-D factorization. We also developed a lifting/ladder structure for the 2-D NSFB. This structure, when coupled with the filters designed via mapping, provides a very efficient implementation that under some additional conditions can be reduced to 1-D filtering operations. Applications of our proposed transform in image denoising and enhancement were studied. In denoising, we studied the performance of the NSCT when coupled with a hard thresholding estimator and a local adaptive shrinkage. For hard thresholding, our results indicate that the NSCT provides better performance than competing transform such as the NSWT and curvelets. Concurrently, our local adaptive shrinkage results are competitive to other denoising methods. In particular, our results show that
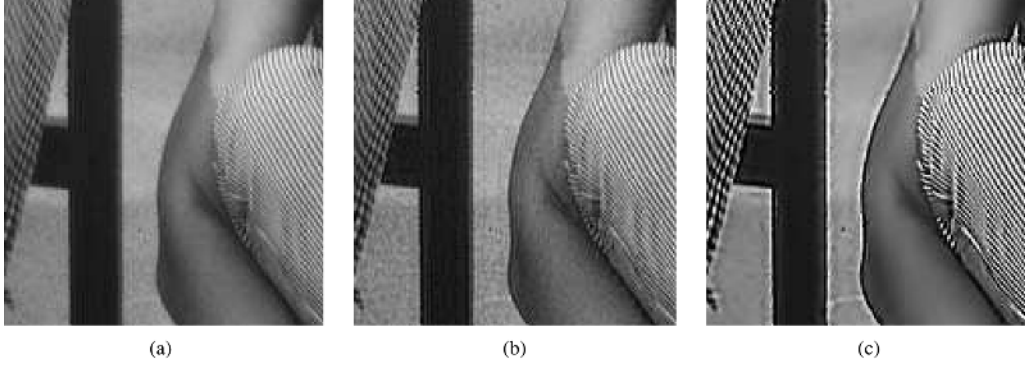
Fig. 12. Image enhancement with the NSCT. (a) Original zoomed "Barbara" image. (b) Enhanced by the NSWT. (c) Enhanced by the NSCT.

a fairly simple estimator in the NSCT domain yields comparable performance to state-of-the-art denoising methods that are more sophisticated and complex. In image enhancement, the results obtained with the NSCT are superior to those of the NSWT both visually and with respect to objective measurements. A MATLAB toolbox that implements the NSCT can be downloaded from MATLAB Central (http://www.mathworks.com/matlabcentral/).

APPENDIX

A. Proof of Proposition 1

Consider the pyramid shown in Fig. 2(a). If $J = 1$, then we have that $\|y_0\|^2 + \|y_1\|^2 \le B_p\|x\|^2$. Now, suppose we have J levels and assume $\sum_{j=0}^{J}\|y_i\|^2 \le B_p^J\|x\|^2$. Then if we further split y_J into y_J' and y_{J+1}, noting that $B_p \ge 1$, we have

$$\sum_{j=0}^{J-1}\|y_j\|^2 + \|y_J'\|^2 + \|y_{J+1}\|^2$$
$$\le \sum_{j=0}^{J-1}\|y_j\|^2 + B_p\|y_J\|^2$$
$$\le B_p\left(\sum_{j=0}^{J-1}\|y_j\|^2 + \|y_J\|^2\right)$$
$$\le B_p^{J+1}\|x\|^2.$$

Thus, by induction, we conclude that $\sum_{j=0}^{J}\|y_i\|^2 \le B_p^J\|x\|^2$ for any $J \ge 1$. A similar argument shows that in the NSDFB with l_j stages, one has that $\sum_{k=0}^{2^{l_j}-1}\|y_{j,k}\|^2 \le B_q^{l_j}\|y_j\|^2$ so that

$$\|y_J\|^2 + \sum_{j=0}^{J-1}\sum_{k=0}^{2^{l_j}-1}\|y_{j,k}\|^2$$
$$\le \|y_J\|^2 + \sum_{j=0}^{J-1}B_q^{l_j}\|y_j\|^2$$
$$\le \|y_J\|^2 + B_q^{\max\{l_j\}}\sum_{j=0}^{J-1}\|y_j\|^2$$
$$\le B^J B_q^{\max\{l_j\}}\|x\|^2.$$

The bound for A is proved similarly, by just reversing the inequalities. ∎

B. Proof of Proposition 2

We prove the claim for the case in which the zeros of $G^{(1D)}(z)$ are distinct. The proof for the case of repeated roots can be handled similarly.

Denote

$$G^{(1D)}(z) = g_0\prod_{i=1}^{n}(z - z_i), \quad g_0 \in \mathbb{C}. \tag{20}$$

Then, sufficiency follows by direct substitution of (9) in (20).

We prove necessity by induction. Suppose $G^{(1D)}(F(x_1, x_2)) = (x_1 - c)^{N_1}L(x_1, x_2)$ for some polynomial $L(x_1, x_2)$. Note that $G^{(1D)}(F(c, x_2)) = 0$ for all x_2 if and only if for some zero z_j of $G^{(1D)}(z)$, $F(c, x_2) = z_j$ for all x_2. So, it follows:

$$(F(x_1, x_2) - z_j)_{x_1=c} = 0 \text{ for all } x_2 \tag{21}$$

which implies that $F(x_1, x_2) = z_j + (x_1 - c)L_1(x_1, x_2)$ with $L_1(x_1, x_2)$ a polynomial. Suppose

$$F(x_1, x_2) = z_j + (x_1 - c)^{k-1}L_{k-1}(x_1, x_2)$$

where $k \le N_1$. By successively applying the chain rule for differentiation, we get

$$0 = \left.\frac{\partial^{k-1}G^{(1D)}(F(x_1, x_2))}{\partial x_1^{k-1}}\right|_{x_1=c}$$
$$= G^{(1D)\prime}(F(c, x_2))\left.\frac{\partial^{k-1}F(x_1, x_2)}{\partial x_1^{k-1}}\right|_{x_1=c}.$$

Because $F(c, x_2) = z_j$ and the z_js are distinct, we have $G^{(1D)\prime}(F(c, x_2)) \ne 0$ and then

$$\left.\frac{\partial^{k-1}F(x_1, x_2)}{\partial x_1^{k-1}}\right|_{x_1=c} = 0$$
$$\Rightarrow \frac{\partial^{k-1}F(x_1, x_2)}{\partial x_1^{k-1}} = (x_1 - c)\tilde{L}_1(x_1, x_2) \tag{22}$$
$$\Rightarrow \frac{\partial F(x_1, x_2)}{\partial x_1} = (x_1 - c)^{k-1}L_{k-1}(x_1, x_2) \tag{23}$$

where (23) follows by successively integrating (22). Combining (23) with (21) we obtain that $F(x_1, x_2) = z_j + (x_1 - c)^k L_k(x_1, x_2)$. By induction, we conclude that

$$F(x_1, x_2) = z_j + (x_1 - c)^{N_1} L_{N_1}(x_1, x_2).$$

If $G^{(1D)}(F(x_1, x_2)) = (x_2 - d)^{N_2} \tilde{L}(x_1, x_2)$ for some polynomial $\tilde{L}(x_1, x_2)$, then a similar argument shows that

$$F(x_1, x_2) = z_i + (x_2 - d)^{N_2} L_{N_2}(x_1, x_2)$$

where z_i is a zero of $G^{(1D)}(z)$. Thus, $z_i = z_j = F(c, d)$. Therefore

$$(x_1 - c)^{N_1} L_{N_1}(x_1, x_2) = (x_2 - d)^{N_2} L_{N_2}(x_1, x_2)$$

and so

$$L_{N_1}(x_1, x_2) = (x_2 - d)^{N_2} L_F(x_1, x_2)$$

and (9) is established with $N_1' = N_1$ and $N_2' = N_2$. ∎

C. Proof of Proposition 3

The proof follows the same lines as for the 1-D case (see [32, pp. 216]). Using the identity

$$\prod_{k=1}^{\infty} \left(\frac{1 + e^{j2^{-k}\omega}}{2} \right)^N = \left(\frac{1 + e^{j\omega}}{\omega} \right)^N$$

we have

$$\Phi(\omega) = \left(\frac{1 + e^{j\omega_1}}{\omega_1} \right)^{N_1} \left(\frac{1 + e^{j\omega_2}}{\omega_2} \right)^{N_2} \prod_{j=1}^{\infty} R_0(2^{-j}\omega).$$

Because H_0 is continuously differentiable at $\omega = 0$, the same also holds for R_0. Now write R_0 in polar coordinates (r, φ) where $r^2 = |\omega_1|^2 + |\omega_2|^2$. Since R_0 is continuously differentiable, for each φ, $R_0(r, \varphi)$ is a continuously differentiable function of r. Since $R_0(0) = 1$, from the mean value theorem we have that for $\epsilon > 0$, and $0 \leq r < \epsilon$

$$|R(r, \varphi)| \leq 1 + |\partial_r R(\rho, \varphi)| r \leq 1 + Kr$$

where $K := \sup\{|R(\rho, \varphi)| : 0 \leq \rho < \epsilon, \varphi \in [0, 2\pi]\} < \infty$. This gives

$$0 \leq r < \epsilon \Longrightarrow \prod_{j=1}^{\infty} R_0(2^{-j}\omega)$$
$$\leq \prod_{j=1}^{\infty} (1 + K2^{-j}r) \leq e^{K\epsilon}$$

where we have used the inequality $1 + r \leq e^r$. Now choose J so that $2^{J-1}\epsilon \leq r \leq 2^J \epsilon$. We then obtain, for $r > \epsilon$

$$\prod_{j=1}^{\infty} |R_0(2^{-j}\omega)|$$
$$= \prod_{j=1}^{J} |R_0(2^{-j}\omega)| \prod_{j=1}^{\infty} |R_0(2^{-j-J}\omega)|$$
$$\leq B^J e^{K\epsilon} \leq C_1 r^{\log_2 B} e^{K\epsilon}$$

so that for each $\omega \in \mathbb{R}^2$

$$\prod_{j=1}^{\infty} R_0(2^{-j}\omega) \leq e^{K\epsilon} \left(1 + C_1 r^{\log_2 B} \right)$$
$$= e^{K\epsilon} \left(1 + C_1(|\omega_1|^2 + |\omega_2|^2)^{\frac{1}{2} \log_2 B} \right).$$

Now, putting it all together, we obtain

$$|\Phi(\omega)| \leq C_2 \frac{1}{|\omega_1|^{N_1}} \frac{1}{|\omega_2|^{N_2}} \left(1 + C_1(|\omega_1|^2 + |\omega_2|^2)^{\frac{1}{2} \log_2 B} \right)$$
$$\leq C_3 \left(\frac{1}{1 + |\omega_1|} \right)^{N_1 - \log_2 B} \left(\frac{1}{1 + |\omega_2|} \right)^{N_2 - \log_2 B}$$

which completes the proof. ∎

REFERENCES

[1] S. Mallat, *A Wavelet Tour of Sinal Processing*, 2nd ed. New York: Academic, 1999.

[2] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*. Englewood Cliffs, NJ: Prentice-Hall, 1995.

[3] J. G. Daugman, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters," *J. Opt Soc. Amer. A*, vol. 2, no. 7, pp. 1160–1169, Jul. 1985.

[4] 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.

[5] R. R. Coifman and D. L. Donoho, "Translation invariant de-noising," in *Wavelets and Statistics*, A. Antoniadis and G. Oppenheim, Eds. New York: Springer-Verlag, 1995, pp. 125–150.

[6] S. G. Chang, B. Yu, and M. Vetterli, "Adaptive wavelet thresholding for image denoising and compression," *IEEE Trans. Image Process.*, vol. 9, no. 9, pp. 1532–1546, Sep. 2000.

[7] L. Sendur and I. W. Selesnick, "Bivariate shrinkage with local variance estimation," *IEEE Signal Process. Lett.*, vol. 9, no. 12, pp. 438–441, Dec. 2002.

[8] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using scale mixtures of Gaussians in the wavelet domain," *IEEE Trans. Image Process.*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.

[9] M. J. Shensa, "The discrete wavelet transform: Wedding the à trous and Mallat algorithms," *IEEE Trans. Signal Process.*, vol. 40, no. 10, pp. 2464–2482, Oct. 1992.

[10] D. L. Donoho, "Wedgelets: Nearly minimax estimation of edges," *Ann. Statist.*, vol. 27, no. 3, pp. 859–897, 1999.

[11] M. B. Wakin, J. K. Romberg, H. Choi, and R. G. Baraniuk, "Wavelet-domain approximation and compression of piecewise smooth images," *IEEE Trans. Image Process.*, vol. 15, no. 5, pp. 1071–1087, May 2006.

[12] E. L. Pennec and S. Mallat, "Sparse geometric image representation with bandelets," *IEEE Trans. Image Process.*, vol. 14, no. 4, pp. 423–438, Apr. 2005.

[13] E. J. Candès and D. L. Donoho, "New tight frames of curvelets and optimal representations of objects with piecewise C'^2 singularities," *Commun. Pure Appl. Math*, vol. 57, no. 2, pp. 219–266, Feb. 2004.

[14] 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.

[15] J. G. Rosiles and M. J. T. Smith, "A low complexity overcomplete directional image pyramid," in *Proc. Int. Conf. Image Processing (ICIP)*, Barcelona, Spain, 2003, vol. 1, pp. 1049–1052.

[16] P. J. Burt and E. H. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.*, vol. 31, no. 4, pp. 532–540, Apr. 1983.

[17] M. N. Do and M. Vetterli, "Framing pyramids," *IEEE Trans. Signal Process.*, vol. 51, no. 9, pp. 2329–2342, Sep. 2003.

[18] 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.

[19] J. L. Starck, F. Murtagh, and A. Bijaoui, *Image Processing and Data Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1998.

[20] M. A. U. Khan, M. K. Khan, and M. A. Khan, "Coronary angiogram image enhancement using decimation-free directional filter banks," in *Proc. Int. Conf. Acoustics, Speech, and Signal Proc. (ICASSP)*, Montreal, QC, Canada, 2004, pp. 441–444.

[21] Z. Cvetkovic and M. Vetterli, "Oversampled filler banks," *IEEE Trans. Signal Process.*, vol. 46, no. 5, pp. 1245–1255, May 1998.

[22] P. P. Vaidyanathan, *Multirate Systems and Filterbanks*. Englewood Cliffs, NJ: Prentice-Hall, 1993.

[23] J. Kovačević and M. Vetterli, "Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for $\mathbb{R}^n$," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 533–555, Mar. 1992.

[24] H. Bölcskei, F. Hlawatsch, and H. G. Feichtinger, "Frame-theoretic analysis of oversampled filter banks," *IEEE Trans. Signal Process.*, vol. 46, no. 12, pp. 3256–3268, Dec. 1998.

[25] J. H. McClellan, "The design of two-dimensional digital filters by transformation," in *Proc. 7th Annu. Princeton Conf. Information Sciences and Systems*, 1973.

[26] S.-M. Phoong, C. W. Kim, P. P. Vaidyanathan, and R. Ansari, "A new class of two-channel biorthogonal filter banks and wavelet bases," *IEEE Trans. Signal Process.*, vol. 43, no. 3, pp. 649–661, Mar. 1995.

[27] D. B. H. Tay and N. G. Kingsbury, "Flexible design of multidimensional perfect reconstruction FIR 2-band filters using transformation of variables," *IEEE Trans. Image Process.*, vol. 2, no. 4, pp. 466–480, Oct. 1993.

[28] A. L. Cunha and M. N. Do, "On two-channel filter banks with directional vanishing moments," *IEEE Trans. Image Process.*, submitted for publication.

[29] S. Mitra and R. Sherwood, "Digital ladder networks," *IEEE Trans. Audio Electroacoust.*, vol. AU-21, no. 1, pp. 30–36, Feb. 1973.

[30] W. Sweldens, "The lifting scheme: A custom-design construction of biorthogonal wavelets," *Appl. Comput. Harmon. Anal.*, vol. 3, no. 2, pp. 186–200, 1996.

[31] R. Jia, "Approximation properties of multivariate wavelets," *Math. Comput.*, vol. 67, pp. 647–665, 1998.

[32] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA: SIAM, 1992.

[33] T. Cooklev, T. Yosbida, and A. Nishihara, "Maximally flat half-band diamond-shaped FIR filters using the Bernstein polynomial," *IEEE Trans. Circuits Syst. II*, vol. 40, no. 11, pp. 749–751, Nov. 1993.

[34] J.-L. Starck, E. J. Candès, and D. L. Donoho, "The curvelet transform for image denoising," *IEEE Trans. Image Process.*, vol. 11, no. 6, pp. 670–684, Jun. 2002.

[35] D. D.-Y. Po and M. N. Do, "Directional multiscale modeling of images using the contourlet transform," *IEEE Trans. Image Process.*, submitted for publication.

[36] S. G. Chang, B. Yu, and M. Vetterli, "Spatially adaptive wavelet thresholding with context modeling for image denoising," *IEEE Trans. Image Process.*, vol. 9, no. 9, pp. 1522–1531, Sep. 2000.

[37] K. V. Velde, "Multi-scale color image enhancement," in *Proc. IEEE Int. Conf. Image Processing*, 1999, vol. 3, pp. 584–587.

[38] G. Ramponi, N. Strobel, S. K. Mitra, and T.-H. Yu, "Nonlinear unsharp masking methods for image contrast enhancement," *J. Electron. Imag.*, vol. 5, no. 3, pp. 353–366, Mar. 1996.

Arthur L. da Cunha received the B.S. degree in electrical engineering from the University of Brasilia, Brasilia, Brazil, in 2000, and the M.S. degree in electrical engineering form the Pontifical Catholic University of Rio de Janeiro (PUC-RIO), Rio de Janiero, Brazil, in 2002. He is currently pursuing the Ph.D. degree in the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana.

During the fall of 2002, he was a graduate student at Georgia Institute of Technology, Atlanta. He is currently a Research Assistant with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign. In the summer of 2005, he was a Research Engineer at NTT DoCoMo Labs USA, and subsequently worked in the Audiovisual Communication Lab, Swiss Federal Institute of Technology, Lausanne, Switzerland. His research interests include signal representation, estimation, and coding.

Mr. da Cunha has won several awards and grants from the Brazilian National Research Council, including two graduate fellowships. In 2005, he won a Best Student Paper Award at the ICASSP Conference (Philadelphia, PA).
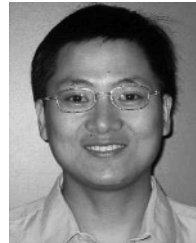
Jianping Zhou (S'05–M'05) received the B.Sc. degree in applied mathematics from Xi'an Jiaotong University, Xi'an, China, in 1997, the M.Eng. degree in signal and information processing from Peking University, Beijing, China, in 2000, and the Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign, Urbana, in 2005.

He is a Member of Technical Staff with the DSP Solution R&D Center, Texas Instruments, Dallas, TX. From 2002 to 2005, he was a Research Assistant with the Coordinated Science Laboratory, University of Illinois. From 2000 to 2002, he was an Assistant Researcher with the Internet Media Group and the Wireless and Networking Group, Microsoft Research Asia, Beijing. From 1997 to 2000, he was a Research Assistant with the National Laboratory on Machine Perception, Beijing, China. His research interests include wavelets, computer vision, digital color processing, image and multidmensional signal processing, multimedia coding and communication, and multimedia delivery over wireless Internet.

Dr. Zhou received an exceptional class scholarship and the "Model of Excellent Students," the highest honor for students in Peking University in 2000, and an exceptional class scholarship and "Model of Excellent Students," the highest honor for students in Xi'an Jiaotong University, in 1995 and 1996.

Minh N. Do (M'02) was born in Thanh Hoa, Vietnam, in 1974. He received the B.Eng. degree in computer engineering from the University of Canberra, Canberra, ACT, Australia, in 1997, and the Dr.Sci. degree in communication systems from the Swiss Federal Institute of Technology Lausanne (EPFL), Lausanne, 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, both at the 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 EPFL in 2001, and a CAREER award from the National Science Foundation in 2003.