New Image Compression Techniques Using Multiwavelets and Multiwavelet Packets

Michael B. Martin and Amy E. Bell, *Member, IEEE*

***Abstract*—Advances in wavelet transforms and quantization methods have produced algorithms capable of surpassing the existing image compression standards like the Joint Photographic Experts Group (JPEG) algorithm. For best performance in image compression, wavelet transforms require filters that combine a number of desirable properties, such as orthogonality and symmetry. However, the design possibilities for wavelets are limited because they cannot simultaneously possess all of the desirable properties. The relatively new field of multiwavelets shows promise in obviating some of the limitations of wavelets. Multiwavelets offer more design options and are able to combine several desirable transform features. The few previously published results of multiwavelet-based image compression have mostly fallen short of the performance enjoyed by the current wavelet algorithms. This paper presents new multiwavelet transform and quantization methods and introduces multiwavelet packets. Extensive experimental results demonstrate that our techniques exhibit performance equal to, or in several cases superior to, the current wavelet filters.**

***Index Terms*—Image compression, multiwavelet packets, multiwavelets.**

I. Introduction

ALGORITHMS based on wavelets[1] have been shown to work well in image compression. Theoretically, multiwavelets should perform even better due to the extra freedom in the design of multifilters. But previously published results still favor wavelets since the effective application of multiwavelets requires solving additional problems to those encountered with wavelets [1], [2]. Theoretical and experimental results in the study of multiwavelets have been steadily progressing and all of the key components for the application of multiwavelets to image compression are now in place. In particular, there now exist methods for: the construction of orthogonal and biorthogonal multifilters with desirable filter properties [3], [4]; good preprocessing techniques [2], [5]; and a method for symmetric signal extension for symmetric–antisymmetric (SA) multiwavelets [3].

Another way to achieve improved compression results over wavelets is to use wavelet packets. Wavelet packets demonstrate a significant improvement in reconstructed image quality over the octave-band wavelet decomposition for some images. This benefit comes from the ability of the wavelet packets to better represent high-frequency content and high-frequency oscillating signals in particular. This allows wavelet packets to perform significantly better than wavelets for the compression of images with a large amount of texture—such as the commonly used Barbara image. Experiments show that wavelet packet techniques applied to such images can outperform wavelet techniques [6], [7]. Researchers also point out that the perceived image quality is significantly improved using wavelet packets instead of wavelets, especially in the textured regions of the images. We introduce a new approach to improving wavelet packet-based image compression: multiwavelet packets.

This paper begins with a brief overview of image compression schemes based on multiwavelets and wavelet packets. Two new techniques for improving the decomposition iteration and zerotree-based quantization for multiwavelets are then presented. Next, multiwavelet packets [8] are defined in terms of multiwavelets and wavelet packets. Extensive experimental results are then presented using recently-constructed orthogonal and biorthogonal (SA) multiwavelets; they illustrate that our new methods improve results over the existing methods (using the best known scalar wavelets) for many test images. Finally, conclusions about the effectiveness and the limitations of these new methods are discussed.

Manuscript received February 13, 2000; revised November 30, 2000. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Eric L. Miller.

M. B. Martin is with Vision III Imaging, Herndon, VA 20170 USA.

A. E. Bell is with the Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24061-0111 USA (e-mail: abell@vt.edu).

Publisher Item Identifier S 1057-7149(01)01654-2.

[1]To distinguish them from multiwavelets, wavelets may also be referred to as *scalar wavelets*.

II. Background

A. Multiwavelets

The wavelet transform is a type of signal transform that is commonly used in image compression. A newer alternative to the wavelet transform is the multiwavelet transform. Multiwavelets are very similar to wavelets but have some important differences. In particular, whereas wavelets have an associated scaling function $\phi(t)$ and wavelet function $\psi(t)$, multiwavelets have two or more scaling and wavelet functions. For notational convenience, the set of scaling functions can be written using the vector notation $\Phi(t) \equiv [\phi_1(t)\ \phi_2(t)\ \cdots\ \phi_r(t)]^T$, where $\Phi(t)$ is called the multiscaling function. Likewise, the multiwavelet function is defined from the set of wavelet functions as $\Psi(t) \equiv [\psi_1(t)\ \psi_2(t)\ \cdots\ \psi_r(t)]^T$. When $r = 1$, $\Psi(t)$ is called a *scalar* wavelet, or simply wavelet. While in principle r can be arbitrarily large, the multiwavelets studied to date are primarily for $r = 2$.

The multiwavelet two-scale equations resemble those for scalar wavelets

$$\Phi(t) = \sqrt{2} \sum_{k=-\infty}^{\infty} H_k \Phi(2t - k), \tag{1}$$

$$\Psi(t) = \sqrt{2} \sum_{k=-\infty}^{\infty} G_k \Phi(2t-k). \tag{2}$$

Note, however, that $\{H_k\}$ and $\{G_k\}$ are *matrix* filters, i.e., H_k and G_k are $r \times r$ matrices for each integer k. The matrix elements in these filters provide more degrees of freedom than a traditional scalar wavelet. These extra degrees of freedom can be used to incorporate useful properties into the multiwavelet filters, such as orthogonality, symmetry, and high order of approximation. The key, then, is to figure out how to make the best use of these extra degrees of freedom. Multifilter construction methods are already being developed to exploit them [3], [4]. However, the multi-channel nature of multiwavelets also means that the subband structure resulting from passing a signal through a multifilter bank is different. Sufficiently different, in fact, so that established quantization methods do not perform as well with multiwavelets as they do with wavelets. In the next section, these key differences are examined and new methods for improving performance are developed.

B. Wavelet Packets

Multiwavelets provide one alternative to the wavelet transform. Another alternative is the wavelet packet transform. Despite its general success, the wavelet transform often fails to accurately capture high-frequency information, especially at low bit rates where such information is lost in quantization noise. Coifman *et al*. developed a technique called wavelet packets that is better able to represent high-frequency information [9].

A multilevel wavelet filter bank involves iterating the lowpass–highpass filtering and downsampling procedure only on the output of the lowpass branch of the previous stage. Coifman *et al.* formulated an extension of the octave-band wavelet decomposition to a full tree decomposition by allowing the lowpass–highpass filtering and downsampling procedure to be iterated also on highpass (bandpass) branches in the tree [9]. They defined the new basis functions, called wavelet packets, as follows.

Let $\phi(t)$ and $\psi(t)$ be the scaling and wavelet functions, respectively, which obey the two-scale equations

$$\phi(t) = \sqrt{2} \sum_{k=-\infty}^{\infty} h_k \phi(2t-k), \tag{3}$$

$$\psi(t) = \sqrt{2} \sum_{k=-\infty}^{\infty} g_k \phi(2t-k). \tag{4}$$

Note that the sequences $\{h_k\}$ and $\{g_k\}$ are the scaling and wavelet filter coefficients. Now let $u_0(t) \equiv \phi(t)$ and $u_1(t) \equiv \psi(t)$, and define

$$u_{2n}(t) = \sqrt{2} \sum_{k=-\infty}^{\infty} h_k u_n(2t-k), \tag{5}$$

$$u_{2n+1}(t) = \sqrt{2} \sum_{k=-\infty}^{\infty} g_k u_n(2t-k). \tag{6}$$

Taking dyadic rescalings and translations of these functions yields a library of functions $\{2^{-j/2} u_n(2^{-j}t-k)\}$. This library is overcomplete, but a proper complete basis can be found by selecting a subset of the library with the right set of parameters $\{n, j, k\}$ [9].

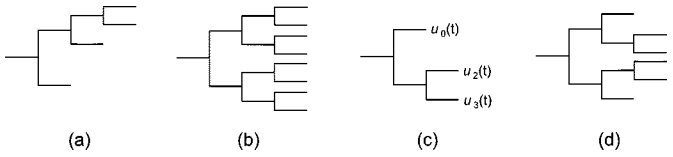

Fig. 1. Possible wavelet packet filter banks: (a) standard wavelet decomposition, (b) full tree (Walsh basis), and (c) and (d) are other possibilities.

This selection of a basis can be viewed in terms of a tree structure, in which the set of elements of each basis corresponds in a one-to-one fashion to a particular set of terminal nodes of a binary tree. Some examples of possible basis selections are shown as trees in Fig. 1. For example, the tree in Fig. 1(a) corresponds to the wavelet octave-band decomposition.

Wavelet packets impose increased computational complexity due to the basis selection process. Selection of a "best" basis for any particular image may be performed in a number of ways. Coifman *et al.* suggested the use of an additive cost function that is applied to each set of parent and child nodes in the pruning process. If the sum of the costs of the children is greater than the parent's cost, the children are pruned; otherwise the children are kept. The performance of this method depends entirely on the choice of cost functions. Some cost functions that have been proposed include: Shannon entropy [10], the number of coefficients in the node that are significant compared to (i.e., greater than) some threshold[2] [6], and the number of bits required to represent all the coefficients in the node (introduced in this paper).

Newer methods for selecting a basis approach the problem from a rate-distortion perspective. Ramchandran and Vetterli proposed a method that attempts to select the set of terminal nodes that are optimal in a rate-distortion sense [11]. Their approach involves the minimization at each branch of a Lagrangian "cost function," $J(\lambda) = D + \lambda R$, where D is the average distortion and R is the target average bit rate. The value of λ that minimizes $J(\lambda)$ determines whether to prune and also gives the best quantizer for that node (which is then used for uniform quantization of the coefficients of that node). More recently, Xiong *et al.* have taken this idea and merged the basis optimization with their space-frequency quantization (SFQ) approach, yielding impressive results [7], [12].

III. New Methods for Multiwavelets

A. Iteration of Decomposition

During a single level of decomposition using a scalar wavelet transform, the two-dimensional (2-D) image data is replaced with four blocks corresponding to the subbands representing either lowpass or highpass filtering in each direction. These subbands are illustrated in Fig. 2(a); for example, the data in subband LH was obtained from highpass filtering of the rows and then lowpass filtering of the columns. The multiwavelets used here have two channels, so there will be two sets of scaling

[2]Usually this threshold is taken to be on the order of the quantization step size.

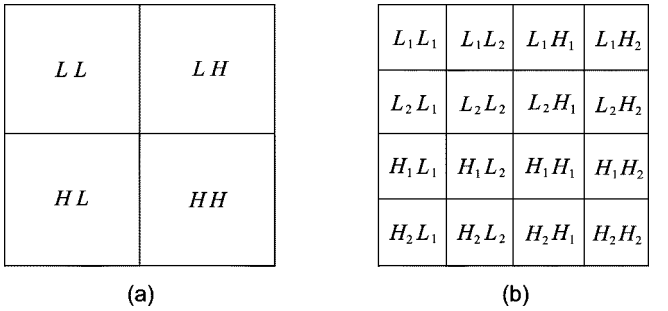

Fig. 2. Image subbands after a single-level decomposition for (a) scalar wavelets and (b) multiwavelets.

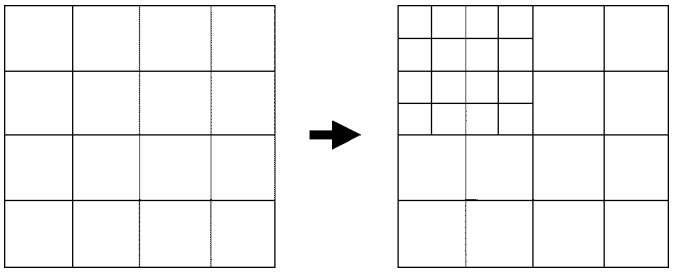

Fig. 3. Conventional iteration of multiwavelet decomposition.

TABLE I
PSNR RESULTS (IN dB) COMPARING STANDARD AND NEW MULTIWAVELET DECOMPOSITION METHODS

Image	Multifilter	Comp. ratio	Decomposition iterated on: L_iL_j Subband (Standard)	Decomposition iterated on: L_1L_1 Subband (New)
Lena	SA4	16:1	33.50	34.66
	BSA9/7	16:1	33.29	34.96
	SA4	32:1	29.85	31.20
	BSA9/7	32:1	29.88	31.94
Barbara	SA4	16:1	28.82	29.58
	BSA9/7	16:1	28.79	30.25
	SA4	32:1	25.78	26.30
	BSA9/7	32:1	25.60	26.80

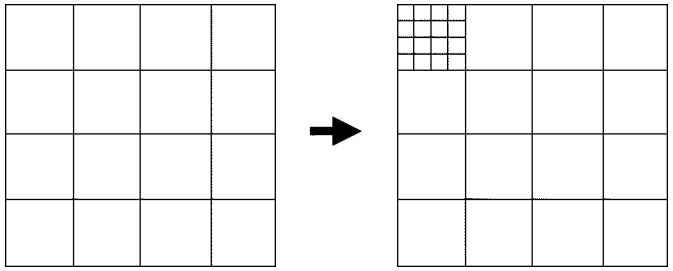

Fig. 4. Proposed iteration method for multiwavelet decomposition. Compare to Fig. 3.

coefficients and two sets of wavelet coefficients. The multiwavelet decomposition subbands are shown in Fig. 2(b). For multiwavelets, the L and H labels have subscripts denoting the channel to which the data corresponds. For example, the subband labeled L_1H_2 corresponds to data from the second channel highpass filter in the horizontal direction and the first channel lowpass filter in the vertical direction.

Scalar wavelet transforms give a single quarter-sized lowpass subband from the original larger subband, as seen in subband LL in Fig. 2(a). In previous multiwavelet literature, multi-

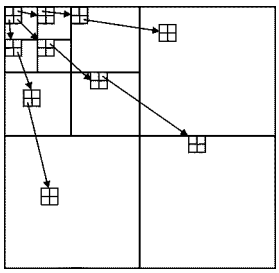

Fig. 5. Illustration of the parent–child relationship in a three-level iterated wavelet decomposition.

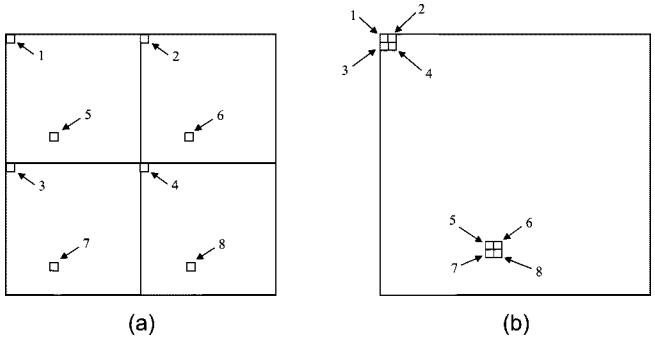

Fig. 6. Illustration of the coefficient shuffling method for one 2×2 subband block in the multiwavelet transform. Selected pixels are numbered to indicate correspondence (a) before shuffling and (b) after shuffling.

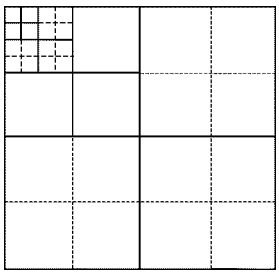

Fig. 7. Subbands in a 2-level multiwavelet decomposition after coefficient shuffling. Solid lines denote new subband boundaries and dashed lines show subband boundaries that are removed by coefficient shuffling.

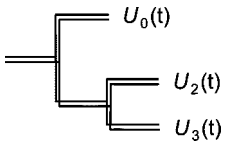

Fig. 8. Possible multiwavelet packet filter bank. Compare to Fig. 1(c).

level decompositions are performed in the same way. The multiwavelet decompositions iterate on the lowpass coefficients from the previous decomposition, [the L_iL_j subbands in Fig. 2(b)], as shown in Fig. 3. In the case of scalar wavelets, the lowpass quarter image is a single subband. But when the multiwavelet transform is used, the quarter image of "lowpass" coefficients is actually a 2×2 block of subbands—one lowpass and three bandpass. This is due to the use of SA multifilters. The $\phi_2(t)$ scaling function associated with the second channel, L_2, is bandpass since its antisymmetric form gives a zero at $z = 1$.

Two conclusions may be drawn from these observations. First, since these four L_iL_j subbands possess different statistical characteristics, mixing them together using the standard

TABLE II
LISTING OF TEST IMAGES AND THEIR SOURCES OF ORIGIN

Name	Source of Origin
Barbara	http://www.icsl.ucla.edu/~ipl/psnr_images.html
Finger	http://noodle.med.yale.edu/~meyer/profile.html
Goldhill	http://www.icsl.ucla.edu/~ipl/psnr_images.html
Gray21	http://sipi.usc.edu/services/database/Database.html
IC	MATLAB Image Processing Toolbox
Lena	http://www.icsl.ucla.edu/~ipl/psnr_images.html
Lighthouse	http://noodle.med.yale.edu/~meyer/profile.html
Mandrill	http://www.code.ucsd.edu/~sherwood/image_examples/chan_coded/chan_coded.html
Testpat_1k	http://sipi.usc.edu/services/database/Database.html
Testpat2	MATLAB Image Processing Toolbox
Yogi	http://saigon.ece.wisc.edu/~waveweb/QMF/software.html

multiwavelet decomposition results in subsequent subbands with mixed data characteristics. This implies that typical quantization schemes that assume the statistics in each subband are either lowpass or highpass will not give the best possible results. Second, since only the L_1L_1 subband is entirely comprised of lowpass characteristics, we only need to perform further iterations on that one subband. Experimental results demonstrate that iterating only on the L_1L_1 subband at each stage in the decomposition does indeed yield better compression than iterating on the four L_iL_j subbands. Some of these results are depicted in Table I. The 1–2 dB performance improvement indicated in Table I is typical of this new decomposition scheme.[3] It is also worth noting the computational savings realized on all iterations subsequent to the first: iterating on only the L_1L_1 subband requires one-fourth the number of computations as iterating over the four L_iL_j subbands. The structure of this new improved multiwavelet decomposition method is illustrated in Fig. 4.

B. Quantization: Shuffling

The quantization method used to generate the results in this paper is the SPIHT zerotree quantizer developed by Said and Pearlman [13]. SPIHT and other zerotree quantizers achieve good performance by exploiting the spatial dependencies of pixels in different subbands of a scalar wavelet transform. It has been noted [14] that there exists a spatial dependence between pixels in different subbands in the form of a child–parent relationship. In particular, each pixel in a smaller subband has four children in the next larger subband in the form of a 2×2 block of adjacent pixels. This relationship is illustrated in Fig. 5, which shows a three-level scalar wavelet decomposition and some sample pixel relations. In this figure, each small square represents a pixel and each arrow points from a particular parent pixel to its 2×2 group of children. The importance of the parent–child relation in quantization is this: if the parent coefficient has a small value, then the children will most likely also have small values; conversely, if the parent has a large value, one or more of the children might also.

The assumptions that the SPIHT quantizer makes about spatial relations between subbands hold well for wavelets, but they do not hold for multiwavelets. More specifically, the three largest highpass subbands in a scalar wavelet transform are each split into a 2×2 block of smaller subbands by the multiwavelet transform, destroying the parent–child relationship that SPIHT presumes [see Fig. 2(b)]. We present a new quantization method that allows multiwavelet decompositions to receive most of the benefits of a SPIHT-like quantizer. The basic idea is to try to restore the spatial features that SPIHT requires for optimal performance. Examination of the coefficients in a single-level multiwavelet transform reveals that there generally exists a large amount of similarity in each of the 2×2 blocks that comprise the L_iH_j, H_iL_j, and H_iH_j subbands, where $i = 1, 2$ and $j = 1, 2$.

This observation suggests the following procedure: rearrange the coefficients in each 2×2 block so that coefficients corresponding to the same spatial locations are placed together. This new procedure will be referred to as *shuffling*. A clearer picture of this is given in Fig. 6. Fig. 6(a) shows one of the 2×2 blocks resulting from a multiwavelet decomposition. Eight pixels (two from each subband) are highlighted and given a unique numeric label. Fig. 6(b) shows the same set of pixels after shuffling. Note that pixels 1–4 map to a 2×2 set of adjacent pixels, as do pixels 5–8. This shuffling procedure restores some of the spatial dependence of the pixels by moving those pixels that correspond to a particular part of the image to the position that they would have been located had a scalar wavelet decomposition been performed.

After shuffling coefficients, a 2-level decomposition iterating on only the L_1L_1 block would look like the one depicted by the solid lines in Fig. 7. Notice that the original subband boundaries—indicated by dashed lines in the figure—have been removed by the shuffling process. Although the coefficients in this multiwavelet decomposition with shuffling scheme are different than the coefficients in a scalar wavelet decomposition, the *structure* of the multiwavelet coefficients in all but the smallest L_iL_i subband is the same as the structure of the coefficients in a four-level scalar wavelet decomposition. In other words, the structure of a four-level scalar wavelet decomposition would be identical to Fig. 7 without the dashed lines—where pixels that correspond to the same spatial locations are next to each other in each subblock (except the smallest L_iL_i subblock). Experimental results in the next section show that this new shuffling scheme improves multiwavelet performance in many cases.

[3]More information about the entries in Table I can be found in Section IV-A.

TABLE III
PSNR RESULTS (IN dB) FOR WAVELETS AND MULTIWAVELETS FOR NATURAL IMAGES (1.0 bpp CORRESPONDS TO AN 8 : 1 COMPRESSION RATIO)

Image	Filter	1.000 bpp	0.500 bpp	0.250 bpp	0.125 bpp
Lena	Bi9/7	38.99	35.62	32.31	29.44
	Bi22/14	**39.23**	**36.06**	**32.79**	**29.87**
	SA4	38.31	34.66	31.20	28.34
	SA4 (sh)	39.06	35.39	31.97	29.08
	ORT4	38.35	34.71	31.24	28.38
	ORT4 (sh)	39.09	35.43	32.01	29.10
	BSA9/7	38.06	34.96	31.94	29.30
	BSA9/7 (sh)	38.49	35.55	32.62	29.81
	BSA7/5	38.49	34.83	31.29	28.38
	BSA7/5 (sh)	39.08	35.43	31.96	28.95
Barbara	Bi9/7	34.58	29.74	26.35	23.81
	Bi22/14	**35.30**	**30.32**	**26.85**	24.00
	SA4	34.60	29.58	26.30	23.82
	SA4 (sh)	34.59	29.50	26.27	23.84
	ORT4	34.66	29.64	26.33	23.86
	ORT4 (sh)	34.65	29.55	26.29	23.83
	BSA9/7	34.71	30.25	26.80	**24.31**
	BSA9/7 (sh)	34.67	30.01	26.60	24.05
	BSA7/5	34.92	29.85	26.48	23.85
	BSA7/5 (sh)	34.91	29.74	26.42	23.74
Goldhill	Bi9/7	35.11	31.78	29.33	27.60
	Bi22/14	35.20	31.86	29.34	**27.74**
	SA4	35.19	31.73	29.08	27.30
	SA4 (sh)	35.30	31.89	29.34	27.54
	ORT4	35.20	31.75	29.10	27.32
	ORT4 (sh)	35.31	31.89	29.35	27.55
	BSA9/7	35.00	31.78	29.27	27.61
	BSA9/7 (sh)	35.03	31.83	**29.46**	27.73
	BSA7/5	35.28	31.82	29.15	27.29
	BSA7/5 (sh)	**35.35**	**31.90**	29.35	27.51
Mandrill	Bi9/7	27.86	24.43	22.30	21.16
	Bi22/14	28.07	24.61	22.35	21.21
	SA4	27.84	24.40	22.24	21.04
	SA4 (sh)	28.15	24.61	22.47	21.16
	ORT4	27.86	24.41	22.25	21.05
	ORT4 (sh)	**28.17**	**24.63**	**22.48**	21.16
	BSA9/7	27.69	24.41	22.22	21.16
	BSA9/7 (sh)	27.97	24.55	22.42	**21.23**
	BSA7/5	27.91	24.49	22.23	21.14
	BSA7/5 (sh)	**28.17**	24.62	22.46	21.20
Finger	Bi9/7	32.73	28.47	24.78	22.36
	Bi22/14	33.92	29.21	25.53	23.13
	SA4	33.26	28.09	24.22	22.10
	SA4 (sh)	33.02	28.08	24.45	22.06
	ORT4	33.38	28.13	24.22	22.12
	ORT4 (sh)	33.13	28.13	24.47	22.07
	BSA9/7	**34.49**	**29.60**	**25.89**	**23.25**
	BSA9/7 (sh)	34.23	29.48	25.78	23.12
	BSA7/5	33.89	28.39	24.21	22.15
	BSA7/5 (sh)	33.63	28.36	24.50	22.08

TABLE IV
PSNR RESULTS (IN dB) FOR WAVELETS AND MULTIWAVELETS FOR SYNTHETIC IMAGES (1.0 bpp CORRESPONDS TO AN 8 : 1 COMPRESSION RATIO)

Image	Filter	1.000 bpp	0.500 bpp	0.250 bpp	0.125 bpp
Gray21	Bi9/7	-	88.17	59.04	49.45
	Bi22/14	-	86.19	57.98	48.61
	SA4	-	∞	67.11	52.44
	SA4 (sh)	-	∞	68.01	**53.54**
	ORT4	-	∞	68.43	52.39
	ORT4 (sh)	-	∞	**69.63**	53.16
	BSA9/7	-	54.62	47.63	45.21
	BSA9/7 (sh)	-	48.61	47.21	44.32
	BSA7/5	-	∞	63.26	50.73
	BSA7/5 (sh)	-	∞	63.30	50.75
Testpat2	Bi9/7	81.52	67.35	62.45	59.44
	Bi22/14	71.55	66.57	61.65	58.79
	SA4	∞	70.12	63.23	60.00
	SA4 (sh)	∞	70.06	62.66	59.40
	ORT4	∞	70.43	**63.28**	**60.02**
	ORT4 (sh)	∞	70.38	62.66	59.41
	BSA9/7	71.55	54.67	46.63	45.51
	BSA9/7 (sh)	48.77	45.94	45.25	45.06
	BSA7/5	∞	**71.73**	63.20	59.79
	BSA7/5 (sh)	∞	71.24	62.37	58.97
Testpat_1k	Bi9/7	46.81	36.25	27.25	22.04
	Bi22/14	48.47	36.67	30.04	24.12
	SA4	52.51	**41.38**	**32.88**	27.14
	SA4 (sh)	53.68	40.15	31.80	26.22
	ORT4	52.25	41.15	32.83	**27.20**
	ORT4 (sh)	**53.89**	39.70	31.96	26.33
	BSA9/7	45.54	35.70	29.18	25.72
	BSA9/7 (sh)	45.09	35.15	28.05	23.46
	BSA7/5	53.05	41.20	32.03	26.98
	BSA7/5 (sh)	51.83	38.61	31.15	25.39
IC	Bi9/7	35.68	30.38	25.85	22.45
	Bi22/14	35.82	30.62	26.03	22.28
	SA4	35.70	30.94	26.10	21.87
	SA4 (sh)	36.38	31.90	**27.03**	22.86
	ORT4	35.71	30.97	26.15	21.91
	ORT4 (sh)	**36.39**	**31.91**	**27.03**	**22.89**
	BSA9/7	35.06	30.30	26.07	22.33
	BSA9/7 (sh)	35.39	30.70	26.45	22.72
	BSA7/5	35.72	30.96	26.38	21.91
	BSA7/5 (sh)	36.25	31.50	26.92	22.58
Yogi	Bi9/7	38.67	29.84	24.84	21.91
	Bi22/14	38.43	29.73	25.00	22.09
	SA4	35.54	28.77	24.13	21.24
	SA4 (sh)	**40.49**	**31.27**	**25.66**	22.38
	ORT4	35.55	28.76	24.15	21.26
	ORT4 (sh)	40.42	31.19	25.65	**22.39**
	BSA9/7	34.97	28.24	24.06	21.50
	BSA9/7 (sh)	37.94	29.80	25.01	22.17
	BSA7/5	35.48	28.68	24.13	21.26
	BSA7/5 (sh)	39.17	30.35	25.12	22.18

C. Multiwavelet Packets

Just as with scalar wavelets, the multiwavelet filter bank procedure involves iterating the filtering operation on the lowpass channel of the filter bank. And, just as with scalar wavelets, new basis functions can be produced by iterating on the highpass channels of multiwavelet filter banks too. This new approach combines wavelet packet decomposition with multiwavelet filters; hence, we call it multiwavelet packet decomposition. We define multiwavelet packets in a manner analogous to the definition of wavelet packets.

Let $U_0(t) \equiv \Phi(t)$ and $U_1(t) \equiv \Psi(t)$, and define

$$U_{2n}(t) = \sqrt{2} \sum_{k=-\infty}^{\infty} H_k U_n(2t-k), \tag{7}$$

$$U_{2n+1}(t) = \sqrt{2} \sum_{k=-\infty}^{\infty} G_k U_n(2t-k). \tag{8}$$

Note the similarity between these mulitwavelet packet equations (7) and (8) and the corresponding wavelet packet equa-

Fig. 9. Barbara at 0.25 bpp (corresponds to a 32 : 1 compression ratio) with Bi22/14 wavelet; PSNR = 26.85 dB.

Fig. 10. Barbara at 0.25 bpp (corresponds to a 32 : 1 compression ratio) with BSA9/7 multiwavelet; PSNR = 26.80 dB.

tions, (5) and (6). Here, the tree structures that represent bases for multiwavelet packets look just like those in Fig. 1 with the exception that the $u_n(t)$ functions in Fig. 1(c) are replaced by the corresponding vector-valued functions $U_n(t)$. For example, the wavelet packet tree in Fig. 1(c) has a multiwavelet version that is shown in Fig. 8.

The basis selection algorithms and cost functions used to prune the resulting tree structure are identical to those of the scalar wavelet packet case with one exception: each branching in the multiwavelet packet tree structure creates four new channels (assuming $r = 2$) instead of just two. Since the multiwavelet packet tree then has four children for each parent, the computational complexity for multiwavelet packets may be higher than for wavelet packets. Cost function based methods will be essentially unaffected because they just operate on all the pixels corresponding to each node; with multiwavelet packets there are four nodes instead of two, but each node represents half as much data. However, methods that perform some form of rate-distortion optimization will require more computation due to the increased number of nodes.

IV. Examples and Further Remarks

A. Multiwavelet Results

Image compression experiments using multiwavelets were conducted both with and without the two new methods: iteration of decomposition and quantization shuffling. Both orthogonal and biorthogonal multiwavelets were tested, and all are from the class of SA multifilters. The orthogonal SA multifilters used are "SA4" and "ORT4" [2], [3]; for biorthogonal SA multifilters we used "BSA7/5" and "BSA9/7" [4]. For comparison, two scalar wavelets were used: the popular biorthogonal "Bi9/7" filter [15] and the recently presented "Bi22/14" biorthogonal filter [16]. All tests used the SPIHT quantizer [13], and no entropy coder was used since we were only interested in comparing the transform and quantization performance. The new L_1L_1 decomposition method was used for all the multiwavelet results.

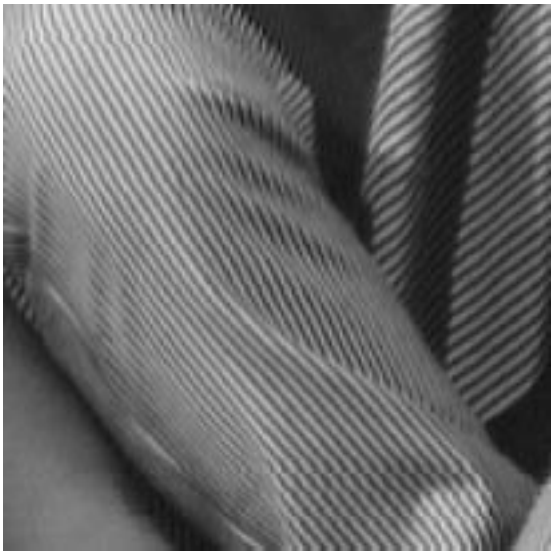

Fig. 11. Original Barbara, showing a close-up of the leg.

The experiments were conducted on many images. The Lena, Barbara, Goldhill, and Mandrill images are the canonical 8 bpp grayscale test images used frequently in the image compression literature (1.0 bpp corresponds to an 8 : 1 compression ratio). The Testpat2 and IC images were taken from the MATLAB Image Processing Toolbox. The remaining images were obtained from various image repositories on the Internet. Table II provides a complete listing of the specific origins for all of the images presented in this paper.

Tables III and IV show PSNR values for reconstructed images. The values shown in boldface represent the best result for each image at each compression level. The "sh" following a filter name in Table III or Table IV indicates that the new shuffling procedure was used in that case. In the following tables, the images are divided into two categories based on their characteristics: "natural" and "synthetic." Natural images are those which derive directly from a real-world source (such as a photograph) and typically have a large amount of low frequency content.

Fig. 12. Lighthouse at 0.25 bpp (corresponds to a 32:1 compression ratio) with Bi22/14 wavelet; PSNR = 26.72 dB.

Fig. 13. Lighthouse at 0.25 bpp (corresponds to a 32:1 compression ratio) with BSA7/5 multiwavelet with shuffling; PSNR = 26.84 dB.

Synthetic images are those which are typically generated by a computer or a similar process; they tend to have more high-frequency content than natural images and they often do not compress as well with traditional image compression transforms like the discrete cosine transform (DCT) and scalar wavelets.

Immediate observations from Tables III and IV suggest that wavelets and multiwavelets have different strengths and weaknesses. Multiwavelets give the best performance on the synthetic images, to the point of achieving lossless compression (i.e., an MSE of zero) on some "geometric" images (Gray21 and Testpat2). Not surprisingly, the short-support orthogonal multifilters SA4 and ORT4 capture the sharp transitions in the synthetic images better than other filters. The longer biorthogonal multifilters BSA9/7 and BSA7/5 perform best on natural images with significant high-frequency content, such as Goldhill, Mandrill, and Finger. Natural images with mostly low-frequency content (e.g., Lena and Barbara) are best compressed with scalar wavelets, where the Bi22/14 filter easily outperforms all other filters. Even for these cases, however, the BSA9/7 multifilter with shuffling sometimes outperforms the commonly-used Bi9/7 filter.

The shuffling procedure adjusts reconstruction performance based on the type and amount of high-frequency content in the image being compressed. Whereas predominantly smooth images like Lena show marked gains when shuffling is used, a corresponding benefit is not realized for images with more high frequency content. In particular, the results for Finger were not improved by shuffling. Similar results are observed with the synthetic images. On the very geometric Testpat2, multiwavelets did better without shuffling, and yet shuffling helped (to varying degrees) on the smoother Gray21, IC, and Yogi images. It should be noted that in those cases where shuffling did improve performance, the improvement could be substantial. For example, in the Lena image, shuffling raised the performance of most multiwavelets from underperforming the Bi9/7 filter to equal or slightly better PSNR levels.

Fig. 14. Original Lighthouse, showing a close-up of the fence and binoculars.

A decrease in performance as a result of shuffling is presumably due to unstructured high-frequency content in the image. Recall that shuffling regroups the pixels so that the original spatial locations are preserved. In low frequency regions, the regrouped pixels are similar; however, in high frequency regions, the regrouped pixels may be quite dissimilar. In this case, the lack of significant spatial dependency means that the benefits derived from the SPIHT quantization are limited. Thus, images like Barbara that lack structure in the bandpass subbands due to its high-frequency content do not experience any benefit from shuffling. However, performance losses due to shuffling are usually small and often occur in cases where the multiwavelets outperform scalar wavelets. For example, shuffling coefficients in the Finger image tended to lower the multiwavelet PSNR results by up to nearly 0.3 dB. Even though the shuffled result for the BSA9/7 multifilter is lower than the unshuffled result at all bit

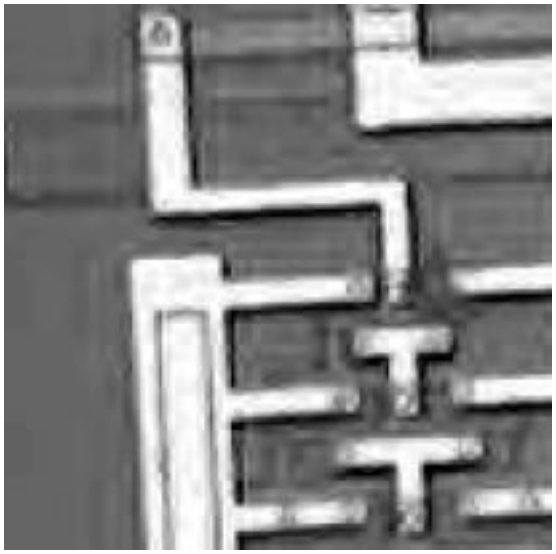

Fig. 15. IC at 0.5 bpp (corresponds to a 16 : 1 compression ratio) with Bi22/14 wavelet; PSNR = 30.62 dB.

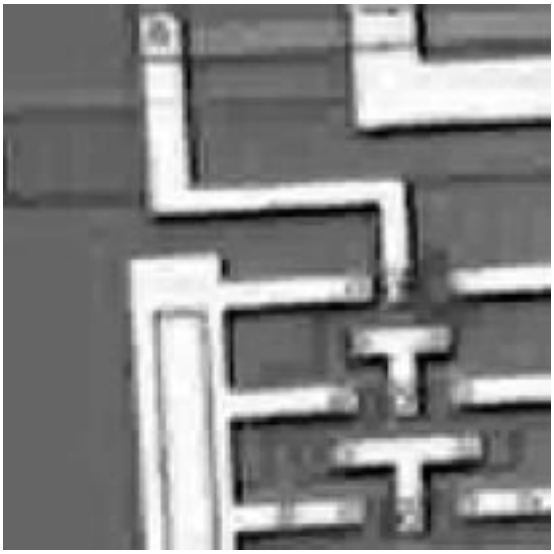

Fig. 16. IC at 0.5 bpp (corresponds to a 16 : 1 compression ratio) with SA4 multifilter with shuffling; PSNR = 31.90 dB.

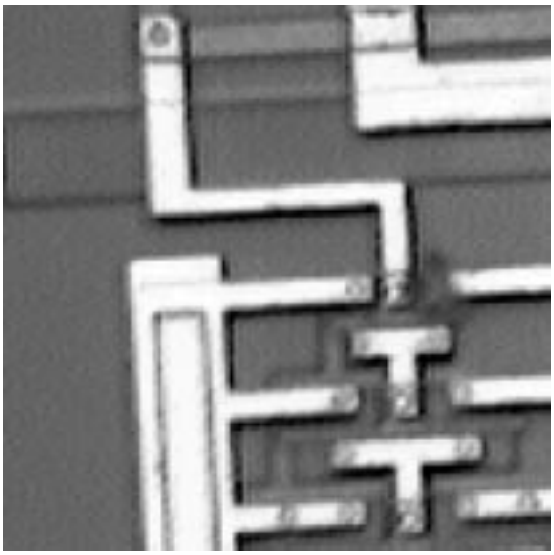

Fig. 17. Original IC image, showing a close-up.

rates, the shuffled result is still at least as good as the Bi22/14 scalar wavelet result at all bit rates. In contrast, when shuffling improves performance, the improvement is often quite significant. Hence, while the type of image being compressed has a significant bearing on whether shuffling would be beneficial, in general it would be safe to use the shuffling method.

In some cases, tests on images that contain large textured regions (like Barbara and Finger) demonstrate that multiwavelets can attain some of the benefits of wavelet packets (by preserving high-frequency patterns that are lost by scalar wavelets with a nonpacket decomposition). An illustration of this feature is depicted in the enlargement of two reconstructions of the Barbara image in Figs. 9 and 10 (the enlargement of the original image is shown in Fig. 11 for comparison). Notice that the pattern in the pants is better preserved by the multiwavelet even though its PSNR is slightly lower than the scalar wavelet PSNR. High-frequency content that is spread over a large image region—or which exhibits oscillations (as in the Barbara image)—is currently best preserved with wavelet packets (as the next section will show), but multiwavelets appear to perform moderately well without the packet-based decomposition.

Two further illustrations visually compare the best wavelet reconstruction with the best multiwavelet reconstruction. Figs. 12 and 13 show two enlargements of reconstructions of Lighthouse—a natural image with a high frequency pattern in the picket fence (the enlargement of the original image is shown in Fig. 14 for comparison). Here the multiwavelet PSNR is slightly higher (0.12 dB) than the wavelet PSNR. Notice in Fig. 13 that the fence details, the edges of the binoculars, and the outline of the background trees is better than in Fig. 12. Figs. 15 and 16 show two enlargements of reconstructions of IC—a synthetic image with large, smooth regions as well as high-frequency edges (the enlargement of the original image is shown in Fig. 17 for comparison). Here the multiwavelet PSNR is significantly higher (1.28 dB) than the wavelet PSNR (shuffling provided 0.96 dB of this improvement). Notice how much closer the multiwavelet with shuffling reconstruction is to the original than the scalar wavelet reconstruction.

The SA4 and ORT4 multiwavelets tend to perform best on synthetic images; it is interesting to note that these two orthogonal multiwavelets show nearly identical performance in most situations. The BSA7/5 multiwavelet performed best on "mixed," natural images like Goldhill and Mandrill. Like the Bi9/7 and Bi22/14 scalar wavelets, the BSA9/7 multiwavelet performed best on natural images. However, while the Bi9/7 and Bi22/14 scalar wavelets perform best on smooth images like Lena, BSA9/7 performs better on images like Finger which have a large amount of structure (and hence some high-frequency patterns) throughout the entire image.

B. Multiwavelet Packet Results

A second set of image compression experiments was conducted using the new multiwavelet packets. Tables V and VI

TABLE V
PSNR RESULTS (IN dB) FOR WAVELET PACKETS AND MULTIWAVELET PACKETS FOR NATURAL IMAGES (1.0 BPP CORRESPONDS TO AN 8:1 COMPRESSION RATIO)

Image	Filter	1.000 bpp	0.500 bpp	0.250 bpp	0.125 bpp
Lena	Bi9/7 (1)	37.84	34.55	31.58	28.78
	Bi9/7 (2)	38.51	35.40	32.21	29.35
	Bi22/14 (1)	37.99	34.93	32.09	29.53
	Bi22/14 (2)	**38.82**	**35.82**	**32.72**	**29.81**
	SA4 (1)	37.82	34.26	30.90	28.13
	SA4 (2)	38.22	34.53	30.99	28.15
	ORT4 (1)	37.83	34.29	30.94	28.18
	ORT4 (2)	38.28	34.57	31.04	28.19
	BSA9/7 (1)	37.36	34.39	31.60	29.20
	BSA9/7 (2)	37.86	34.65	31.69	29.24
	BSA7/5 (1)	38.21	34.65	31.10	28.08
	BSA7/5 (2)	38.44	34.68	31.11	28.14
Barbara	Bi9/7 (1)	35.02	30.67	27.37	24.89
	Bi9/7 (2)	35.84	31.30	27.83	25.18
	Bi22/14 (1)	35.71	31.23	27.85	25.21
	Bi22/14 (2)	**36.42**	**31.84**	**28.30**	**25.50**
	SA4 (1)	34.45	29.52	26.48	24.02
	SA4 (2)	34.34	29.62	26.75	24.37
	ORT4 (1)	34.50	29.59	26.54	24.08
	ORT4 (2)	34.45	29.68	26.78	24.41
	BSA9/7 (1)	33.67	29.26	26.38	24.35
	BSA9/7 (2)	34.61	30.02	26.78	24.24
	BSA7/5 (1)	34.92	29.85	26.58	23.87
	BSA7/5 (2)	34.69	29.92	26.98	24.54
Goldhill	Bi9/7 (1)	34.74	31.60	29.26	27.37
	Bi9/7 (2)	35.07	31.95	29.49	27.57
	Bi22/14 (1)	34.89	31.78	29.40	27.61
	Bi22/14 (2)	35.17	**32.01**	**29.52**	**27.75**
	SA4 (1)	34.88	31.58	29.02	27.13
	SA4 (2)	35.06	31.73	29.09	27.18
	ORT4 (1)	34.75	31.46	29.02	27.16
	ORT4 (2)	35.05	31.75	29.10	27.18
	BSA9/7 (1)	34.13	31.25	28.99	27.16
	BSA9/7 (2)	34.84	31.65	29.25	27.58
	BSA7/5 (1)	34.94	31.64	29.11	27.21
	BSA7/5 (2)	**35.23**	31.82	29.16	27.27
Mandrill	Bi9/7 (1)	26.97	23.98	21.96	20.95
	Bi9/7 (2)	27.53	24.37	22.23	**21.16**
	Bi22/14 (1)	27.08	24.08	22.06	20.97
	Bi22/14 (2)	27.73	**24.45**	**22.28**	21.14
	SA4 (1)	27.21	24.06	21.96	20.94
	SA4 (2)	27.76	24.36	22.21	21.04
	ORT4 (1)	27.24	24.06	21.96	20.95
	ORT4 (2)	27.73	24.32	22.20	21.04
	BSA9/7 (1)	27.27	23.58	21.73	20.75
	BSA9/7 (2)	27.66	24.41	22.22	**21.16**
	BSA7/5 (1)	27.30	24.13	22.03	21.02
	BSA7/5 (2)	**27.84**	24.44	22.21	21.14
Finger	Bi9/7 (1)	34.42	29.39	25.67	22.90
	Bi9/7 (2)	34.59	29.64	25.92	23.26
	Bi22/14 (1)	35.56	30.63	26.78	23.80
	Bi22/14 (2)	**35.63**	**30.65**	**26.77**	**23.81**
	SA4 (1)	33.25	28.08	24.21	22.10
	SA4 (2)	34.10	28.99	25.38	22.85
	ORT4 (1)	33.36	28.12	24.21	22.11
	ORT4 (2)	34.26	29.06	25.41	22.87
	BSA9/7 (1)	34.52	29.55	25.95	23.35
	BSA9/7 (2)	34.56	29.75	26.26	23.57
	BSA7/5 (1)	34.09	28.70	24.83	22.66
	BSA7/5 (2)	34.60	29.27	25.54	22.94

TABLE VI
PSNR RESULTS (IN dB) FOR WAVELET PACKETS AND MULTIWAVELET PACKETS FOR SYNTHETIC IMAGES (1.0 BPP CORRESPONDS TO AN 8:1 COMPRESSION RATIO)

Image	Filter	1.000 bpp	0.500 bpp	0.250 bpp	0.125 bpp
Gray21	Bi9/7 (1)	-	72.64	59.41	49.66
	Bi9/7 (2)	-	71.34	58.32	49.56
	Bi22/14 (1)	-	71.57	57.01	49.41
	Bi22/14 (2)	-	70.31	56.11	49.42
	SA4 (1)	-	**77.50**	65.56	52.19
	SA4 (2)	-	76.90	64.01	**52.27**
	ORT4 (1)	-	77.39	**65.77**	52.01
	ORT4 (2)	-	76.63	63.94	52.11
	BSA9/7 (1)	-	68.28	55.84	48.71
	BSA9/7 (2)	-	68.31	56.92	49.56
	BSA7/5 (1)	-	77.12	62.99	50.75
	BSA7/5 (2)	-	73.92	60.16	50.98
Testpat2	Bi9/7 (1)	81.11	75.12	66.05	61.27
	Bi9/7 (2)	79.67	75.16	66.00	61.43
	Bi22/14 (1)	73.42	69.19	65.04	60.44
	Bi22/14 (2)	73.44	69.28	65.13	60.38
	SA4 (1)	∞	78.23	70.51	63.42
	SA4 (2)	∞	78.44	70.60	63.42
	ORT4 (1)	96.30	77.97	70.66	**63.46**
	ORT4 (2)	96.30	77.54	70.81	**63.46**
	BSA9/7 (1)	87.84	79.86	70.52	62.32
	BSA9/7 (2)	87.84	80.73	70.98	62.69
	BSA7/5 (1)	∞	**82.32**	70.67	62.86
	BSA7/5 (2)	∞	81.82	**71.12**	62.86
Testpat_1k	Bi9/7 (1)	45.86	35.95	29.05	24.42
	Bi9/7 (2)	46.16	36.57	30.45	26.14
	Bi22/14 (1)	47.28	35.99	29.89	24.06
	Bi22/14 (2)	47.05	37.40	30.80	26.68
	SA4 (1)	52.51	**41.38**	32.88	27.14
	SA4 (2)	50.45	40.86	**33.08**	28.19
	ORT4 (1)	52.25	41.15	32.83	27.20
	ORT4 (2)	50.31	40.45	**33.08**	**28.21**
	BSA9/7 (1)	46.04	35.42	29.20	25.63
	BSA9/7 (2)	45.82	35.64	29.40	25.74
	BSA7/5 (1)	**53.05**	41.20	32.03	26.98
	BSA7/5 (2)	50.92	40.90	32.66	27.78
IC	Bi9/7 (1)	34.01	29.82	25.83	22.21
	Bi9/7 (2)	35.09	30.55	26.10	22.29
	Bi22/14 (1)	34.28	29.91	26.07	**22.36**
	Bi22/14 (2)	35.01	30.43	25.94	21.80
	SA4 (1)	34.64	29.97	25.55	21.61
	SA4 (2)	35.47	30.89	26.01	21.63
	ORT4 (1)	34.65	29.99	25.62	21.67
	ORT4 (2)	35.48	30.93	26.06	21.72
	BSA9/7 (1)	33.31	29.49	25.91	22.08
	BSA9/7 (2)	34.86	30.28	26.07	22.25
	BSA7/5 (1)	34.72	29.78	25.74	22.20
	BSA7/5 (2)	**35.72**	**30.95**	**26.37**	21.89
Yogi	Bi9/7 (1)	**38.67**	**29.84**	24.83	21.91
	Bi9/7 (2)	**38.67**	**29.84**	24.83	21.91
	Bi22/14 (1)	38.43	29.73	**25.00**	**22.09**
	Bi22/14 (2)	38.43	29.73	**25.00**	**22.09**
	SA4 (1)	35.54	28.77	24.13	21.23
	SA4 (2)	35.54	28.77	24.13	21.23
	ORT4 (1)	35.55	28.76	24.15	21.25
	ORT4 (2)	35.55	28.76	24.15	21.25
	BSA9/7 (1)	33.36	27.59	23.74	21.31
	BSA9/7 (2)	34.97	28.24	24.05	21.50
	BSA7/5 (1)	35.48	28.68	24.13	21.26
	BSA7/5 (2)	35.48	28.68	24.13	21.26

show PSNR values for the reconstructed images using wavelet packets and multiwavelet packets. The number in parentheses following a filter name in either of these tables indicates which cost function was used for that case. Cost function "1" computes the cost as the number of significant coefficients[4] in the tested node. Cost function "2," a new measure that we are proposing, computes the cost as the total number of bits required in the binary representation of all the coefficients in that node.

The results in Tables V and VI suggest multiwavelet packet performance is mixed. While the multiwavelet packets typically give the best results for the synthetic images, wavelet packets give the best results for the natural images with few exceptions. These results are similar to the nonpacket tests in Tables III and IV, except now the multiwavelets show weaker results for the natural images. For the Barbara image, the best wavelet packet result at each bit rate outperformed the best multiwavelet packet result by between 1 and 1.5 dB. However, the multiwavelet packets performed best at 1 bpp for the Goldhill and Mandrill images, and for Mandrill a multiwavelet packet result could be chosen at each bit rate that essentially equaled the best scalar wavelet packet result. Multiwavelet packets gave predominantly better results on the synthetic images. In particular, the SA4 and BSA7/5 multiwavelet packets achieve perfect reconstruction on the highly geometric Testpat2 image at 1 bpp.

The authors presume that multiwavelet packets performed relatively poorly on the natural images because the iterated multiwavelet transform produces a different subband structure than the wavelet transform. The standard wavelet transform structure assumed by SPIHT does not match the multiwavelet packet structure very well and hence some performance is lost due to nonideal quantization. While the coefficient shuffling method introduced in this paper improved multiwavelet performance with the SPIHT quantizer, no similar method currently exists for multiwavelet packets and SPIHT-like quantizers. The use of a different quantization method, such as a uniform scalar quantizer, should give better results for multiwavelet packets. Cost function "2" gives the best results in most cases; this is to be expected since this cost function was chosen to work well with the SPIHT quantization method used here. However, a basis selection method based on a rate-distortion approach [7], [11], [12] would most likely result in better performance than the simple cost-function method used here.

V. Conclusions

Two new methods for improving the multiwavelet transform have been proposed in this paper: a new multiwavelet decomposition that iterates only on the L_1L_1 subband, and a coefficient shuffling method to improve performance with zerotree-based quantizers. Both methods have been shown to improve the performance of multiwavelet image compression in many cases. While the improved decomposition iteration gives uniformly better results, the performance gains of shuffling depend on the image content. Shuffling helps most images with more low-frequency content; images with more high-frequency content typically realize no significant performance benefit and in some cases, performance is degraded. However, performance decreases tend to be quite small whereas performance increases from shuffling are often quite significant. Although this result was not unexpected, a more thorough understanding of why shuffling improves the compression and reconstruction of images with predominantly low-frequency—and not high-frequency—content remains to be developed.

The other new contribution of this paper—multiwavelet packets—outperformed wavelet packets on images containing large amounts of high-frequency content that is either mostly unstructured (as in Mandrill) or geometric or regular in nature (e.g., Testpat2, Testpat_1k, and IC). However, wavelet packets exhibited better performance on most of the natural images. Moreover, it was shown that multiwavelets can achieve some of the benefits of wavelet packets without the computational expense of the packet-based decomposition.

It should be pointed out that the scalar wavelets used here represent the best known filters published after years of study. In contrast, the multifilters used here are still quite new—many have only been discovered within the past two years. Nevertheless, the multiwavelets used in this paper depicted performance equal to the best scalar wavelets in many cases. While the Bi22/14 scalar wavelet gave consistently good performance for natural images, in most cases, a multiwavelet should give similar performance with lower computational complexity. Similarly, in many cases, a multiwavelet packet resulted in similar performance with lower computational complexity than the best scalar wavelet packets. This indicates that multiwavelets are a viable alternative to scalar wavelets in many situations.

The techniques presented in this paper produce some of the best-reported results to date for multiwavelet-based image compression compared to wavelet-based methods. Nonetheless, there is always room for improvement. Since multiwavelets are a relatively new subject of study, only a few construction methods have been published. While the latest published methods can construct SA multiwavelets with desirable magnitude response characteristics, most current filters have few orders of approximation.[5] Future construction methods that add higher orders of approximation while preserving the desirable features of the current methods would most likely result in multifilters that perform even better in image compression applications. Also, methods for reducing the computational complexity of multiwavelets would be helpful, such as factoring the multifilter into a cascade of shorter multifilters (as Meyer *et al.* do for scalar wavelets [6]) and implementation of the multifilter via the lifting scheme.[6] Finally, good results have been presented for applying multiwavelets to the denoising of 1-D and 2-D signals [1], [5], [17]. Combined with the success shown here for multiwavelet image compression, it seems likely that multiwavelets may work well for the compression of noisy images.

[4]In this case, the threshold used for significance testing is simply 0.5, the threshold below which a coefficient will be converted to 0 during integer conversion.

[5]For example, the SA4 multiwavelet has only one approximation order, while the Bi9/7 scalar wavelet has four.

[6]Use of the lifting method could also result in multifiltering methods which can be performed "in place."

REFERENCES

[1] V. Strela, P. N. Heller, G. Strang, P. Topiwala, and C. Heil, "The application of multiwavelet filter banks to image processing," *IEEE Trans. Image Processing*, vol. 8, pp. 548–562, Apr. 1999.

[2] J. Y. Tham, L.-X. Shen, S. L. Lee, and H. H. Tan, "A general approach for analysis and application of discrete multiwavelet transforms," *IEEE Trans. Signal Processing*, vol. 48, pp. 457–464, Feb. 2000.

[3] T. Xia and Q. Jiang, "Optimal multifilter banks: Design, related symmetric extension transform and application to image compression," *IEEE Trans. Signal Processing*, vol. 47, pp. 1878–1889, July 1999.

[4] S. S. Goh, Q. Jiang, and T. Xia, Construction of biorthogonal multiwavelets using the lifting scheme, preprint, 1998.

[5] V. Strela and A. T. Walden, "Orthogonal and biorthogonal multiwavelets for signal denoising and image compression," *Proc. SPIE*, vol. 3391, pp. 96–107, 1998.

[6] F. G. Meyer, A. Z. Averbuch, and J. O. Strömberg, "Fast adaptive wavelet packet image compression," *IEEE Trans. Image Processing*, vol. 9, pp. 792–800, May 2000.

[7] Z. Xiong, K. Ramchandran, and M. T. Orchard, "Wavelet packet image coding using space-frequency quantization," *IEEE Trans. Image Processing*, vol. 7, pp. 892–898, June 1998.

[8] M. B. Martin and A. E. Bell, "Multiwavelet packet image compression: Theory and results," in *Proc. IEEE DSP Workshop*, Oct. 2000.

[9] R. R. Coifman, Y. Meyer, and M. V. Wickerhauser, "Wavelet analysis and signal processing," in *Wavelets and Their Applications*. Boston, MA: Jones and Bartlett, 1992, pp. 153–178.

[10] R. R. Coifman and M. V. Wickerhauser, "Entropy-based algorithms for best basis selection," *IEEE Trans. Inform. Theory*, vol. 38, pp. 713–718, Mar. 1992.

[11] K. Ramchandran and M. Vetterli, "Best wavelet packet bases in a rate-distortion sense," *IEEE Trans. Image Processing*, vol. 2, pp. 160–175, Apr. 1993.

[12] Z. Xiong, K. Ramchandran, and M. T. Orchard, "Space-frequency quantization for wavelet image coding," *IEEE Trans. Image Processing*, vol. 6, pp. 677–693, May 1997.

[13] A. Said and W. A. Pearlman, "A new, fast, and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 6, pp. 243–250, June 1996.

[14] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. Image Processing*, vol. 41, pp. 3445–3462, Dec. 1993.

[15] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image coding using wavelet transform," *IEEE Trans. Image Processing*, vol. 1, pp. 205–220, Feb. 1992.

[16] D. Wei, H.-T. Pai, and A. C. Bovik, "Antisymmetric biorthogonal coifiets for image coding," in *Proc. IEEE Int. Conf. Image Processing*, Oct. 1998.

[17] V. Strela, "Multiwavelets: Theory and applications," Ph.D. dissertation, Mass. Inst. Technol., Cambridge, 1996.

Michael B. Martin was born in Blacksburg, VA, in 1973. He received the B.S. degree in physics in 1993 from Grinnell College, Grinnell, IA. After completing simultaneous M.S. degrees in math and physics at Virginia Tech, Blacksburg, in 1997, he transferred to the Bradley Department of Electrical and Computer Engineering, Virginia Tech, and received the M.S. degree in electrical engineering in 1999.

His academic interests include image processing, wavelets, filter banks, multirate DSP, adaptive filtering, and video compression. He is currently a Senior Software Engineer with Vision III Imaging, Inc., Herndon, VA, a company specializing in imaging technology.

Amy E. Bell (S'89–M'97) was born in Sewickley, PA. She received the B.S. degree in industrial engineering and the M.S. degree in electrical engineering both from the University of Pittsburgh, Pittsburgh, PA. In 1997, she received the Ph.D. degree in electrical engineering from the University of Michigan, Ann Arbor.

She is presently an Assistant Professor with the Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg. Her current research interests include wavelet-based image compression techniques, wavelet-based modulation algorithms, and data fusion techniques for improved geolocation estimates in wireless networks.

Dr. Bell received a National Science Foundation CAREER award in 1999. She has also received two awards for teaching excellence.