ISO/IEC JTC1/SC29/WG1N1912
December 2000

ISO/IEC JTC1/SC29/WG1
(ITU-T SG8)

Coding of Still Pictures

JBIG	**JPEG**
Joint bi-level Image Experts Group	Joint Photographic Experts Group

TITLE:	Report on Multicomponent Wavelet Decorrelation with Orthogonal Filter Banks
SOURCE :	Chris Brislawn, Michelle Pal, Sue Mniszewski Los Alamos National Lab, USNB
PROJECT :	JPEG 2000
STATUS :	Report
REQUESTED ACTION :	To be presented at WG1 meeting in New Orleans, Dec. 2000
DISTRIBUTION:	WG1 web pages

Contact :
ISO/IEC JTC1/SC29/WG1 Convenor - Dr. Daniel Lee
Hewlett-Packard Company, 1501 Page Mill Road, Palo Alto, CA 94304-1126
Tel: +1 415 857 3946, Fax: +1 415 857 4691, E-mail: dlee@hplabs.hp.com

1 Introduction

Los Alamos National Lab has been investigating decorrelation of multicomponent images using nonlinear phase wavelet filter banks to see whether such capabilities should be actively supported by Part 2 of the JPEG-2000 standard. Since such transforms are linear, they are automatically covered under the provisions in Annex I for user-supplied linear decorrelating transforms. Implementing such filter banks as general linear transforms, however, would require that they be signaled in the codestream via their matrix representation, which creates significantly more codestream overhead than would be required to signal these transforms via their lifting coefficients, as is done with linear phase wavelets.

Accordingly, it would be desirable to make the syntax for signaling user-supplied filter banks general enough to signal nonlinear phase filter banks, provided there is evidence that they may prove useful for some component-decorrelating applications. The lifting factorizations of linear phase filter banks have special symmetry properties that are currently exploited by the signaling in such a way as to preclude signaling nonlinear phase filter banks. Note that nonlinear phase filter banks are **NOT** being proposed for spatial wavelet transformations, which are currently restricted to whole-sample symmetric (WS) and half-sample symmetric (HS) filter banks implemented using either symmetric extension (Annex G) or single-sample overlap extension (Annex H).

Additional motivation for this investigation is the fact that many of the spectral transforms traditionally used in multi- or hyperspectral image processing are orthogonal transforms; e.g., KLT's and the variants of the Fourier transform. Orthogonal transforms are desirable and popular in many scientific and statistical signal processing applications. As is well-known, the only orthogonal, linear phase, FIR two-channel wavelet filter bank is the Haar wavelet, which has limited utility in high-performance image coding systems. Enabling practical orthogonal filter banks for multispectral (or other multicomponent) imaging applications therefore requires the standard to enable effective signaling of such nonlinear phase filter banks via their lifting coefficients.

2 Experimental Design

All experiments have the general form of comparing rate-distortion performance of algorithms using different wavelet filter banks for component decorrelating transforms. We are in the process of investigating performance metrics other than SNR, such as classification error rate in common multispectral pixel classification algorithms, but do not have results to report as of this writing.

Since orthogonal filter banks have nonlinear phase characteristics, they are more difficult (but not impossible) to implement via symmetric extension methods. To keep our initial investigations simple, we have opted to implement nonlinear phase filter banks via circular convolution, although we acknowledge that this places certain limitations on filter bank implementation and performance and may not be optimal for production applications. To compensate somewhat for this restriction, we compare nonlinear phase filter banks implemented with circular convolution against the 9-7 WS filter bank implemented two different ways: using symmetric extension (the default implementation in the JPEG standard at present) and using circular convolution. This approach allows us to make more informed judgements about whether performance comparisons are primarily attributable to intrinsic filter bank characteristics or whether relative performance is dominated by coding behavior at signal boundaries.

For the sake of comparing orthogonal filter bank performance against the 9-7 WS filter bank, we have chosen an orthogonal Daubechies symlet, sym4. This is an 8-tap/8-tap paraunitary filter bank with the same number of vanishing moments (4 analysis, 4 synthesis) as the 9-7 filter bank. Sym4 has not received widespread usage in source coding applications as far as we are aware, and its (spatial) image coding performance is unknown. In contrast, the 9-7 WS filter bank has been studied extensively and is widely regarded as one of the best irreversible filter banks known for image coding. Two other orthogonal filter banks were tested, the coiflet with 2 vanishing moments and the 4-moment minimal phase Daubechies orthogonal wavelet, but, as will be shown in the first set of experimental results, they did not perform quite as well as the 4-moment symlet. Longer-term possibilities include studying the performance of data-dependent (optimized) orthogonal filter banks.

The 9-7 filter bank using symmetric extension was applied using the component mixing module in VM 8.0. To implement both filter banks via circular convolution, the source image cube was filtered in the cross-component direction using Matlab. In order to input the decorrelated image components to VM 8.0 for compression, the floating point Matlab output was first quantized to 24-bit precision using a uniform scalar quantizer scaled to achieve zero overload distortion. Because of significant data outliers, this strategy seems to produce higher reconstructed SNR than quantization with a smaller loading factor that incurs overload distortion. After 24-bit quantization, the decorrelated components were input to the VM and subjected to the default rate-constrained bit allocation, quantization, and compression, with no further component mixing, as shown in

Figure 1.

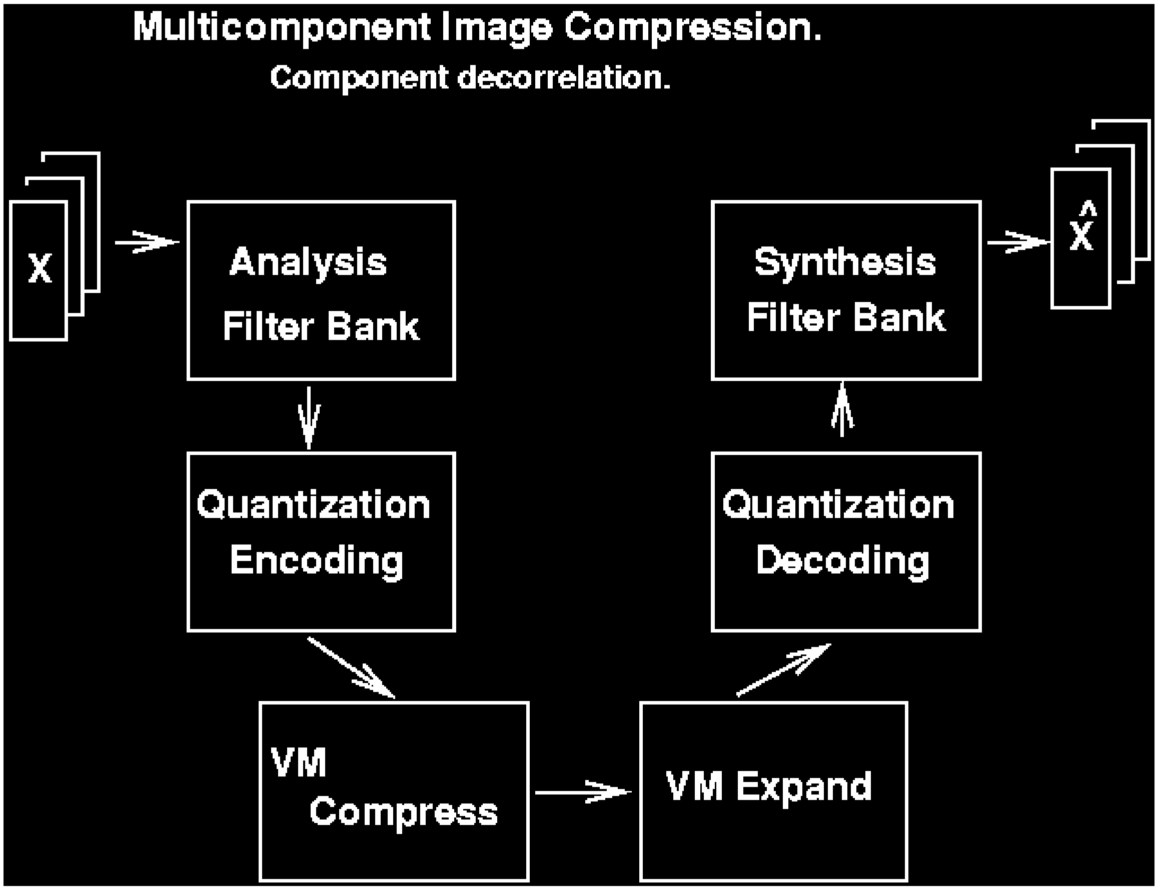

Figure 1. General experimental setup for decorrelation performance comparisons.

3 Experimental Results

We describe several experiments conducted on the AVIRIS Jasper hyperspectral image included in the JPEG-2000 test corpus.

3.1 AVIRIS decorrelation in the image domain

In this experiment, the filter banks were applied to the 224-dimensional spectral vectors of the Jasper image as described above. As a preliminary check on performance, the theoretical coding gain was computed on the decorrelated floating point spectral subbands obtained using several different filter banks, all implemented via circular convolution; the results are shown in Table 1.

Table 1. Coding gain comparison of several filter banks on AVIRIS Jasper data set.

Filter Bank	Coding Gain
9x7 (Biorthogonal 4.4)	29.89
Daubechies N=4	26.55
Symlet 4	32.39
Coiflet 2	30.51

Of the 3 orthogonal filter banks tested, the 4-moment symlet was almost 2 dB better than the 2-moment coiflet and almost 4 dB better than the 4-moment Daubechies minimal-phase wavelet. Note also that the symlet's theoretical coding gain was 2.5 dB higher than the coding gain produced by the 9-7 filter bank.

The rate-distortion characteristics for decorrelation, compression and expansion using these filter banks are shown in Figure 2 and Figure 3. At low rates (Figure 2), several observations stand out. First, the 4-moment minimal phase wavelet consistently lags the other filter banks in performance while the two 9-7 filter bank implementations are consistently the best, a bit over one dB better than the 4-moment minimal phase wavelet. The symmetric extension implementation of the 9-7 (labeled "VM DWT") is slightly better than the circular implementation (labeled "bior4.4"), but only by a couple tenths of a dB. An unknown amount of this already small discrepancy is attributable to the fact that the circular convolution data was quantized to 24 bits just to fit it into the VM, making the intrinsic difference between circular and symmetric extension with the 9-7 filter bank even smaller than it appears in these plots. The difference is certainly far less than would be the case if the comparison were made for spatial (2-dimensional) image filtering. The coiflet and symlet are in the middle, and very close in performance down to 0.25 bpppb, lagging the two 9-7 implementations by about half a dB. At 0.125 bpppb the performance of the coiflet drops to the level of the minimal-phase wavelet while the symlet stays within half a dB of the 9-7 implementations.

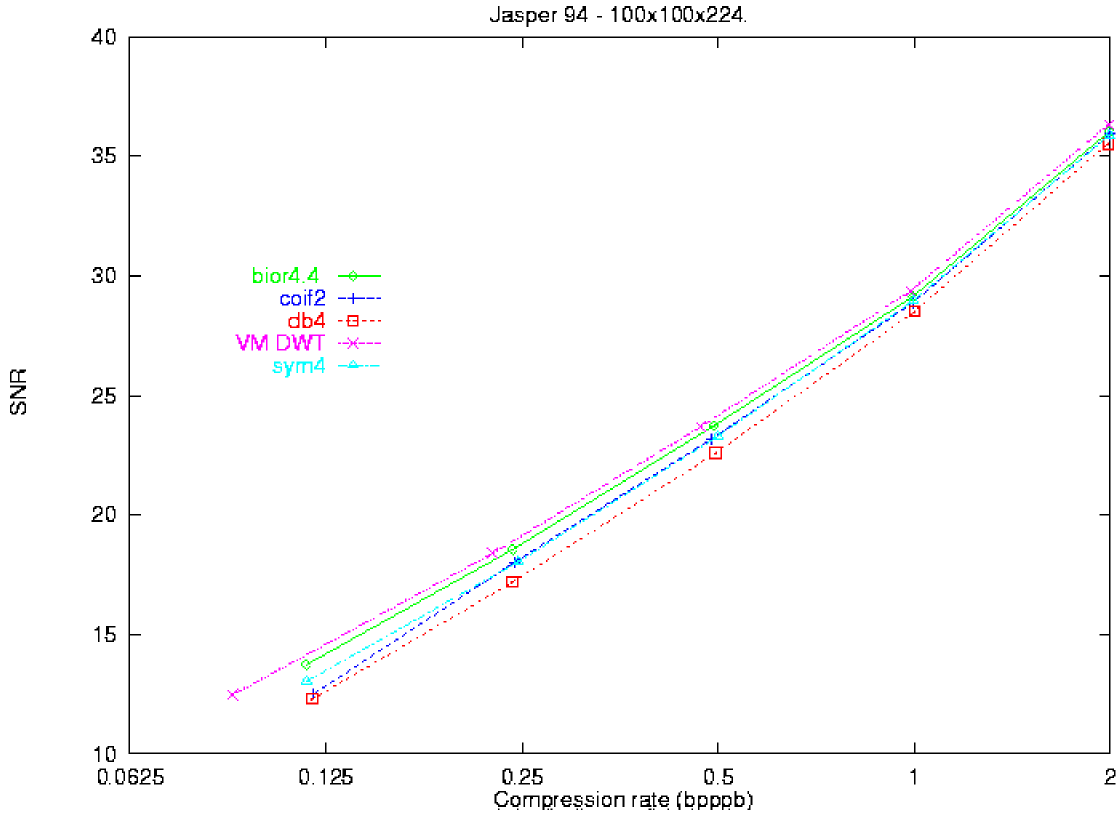

Figure 2. Jasper rate-distortion characteristics at rates below 2 bpppb.

At high rates (Figure 3), the 4-moment minimal phase wavelet continues to be the poorest performer, and the 2 implementations of the 9-7 continue to be almost indistinguishable, but at 4 bpppb the performance of the coiflet and the symlet cross over the 9-7 performance curves. At 8 bpppb the coiflet is a few tenths of a dB better than the 9-7 options and the symlet about one dB better. This finding is noteworthy since hyperspectral imaging applications will probably require coding at very high (even lossless) rates.

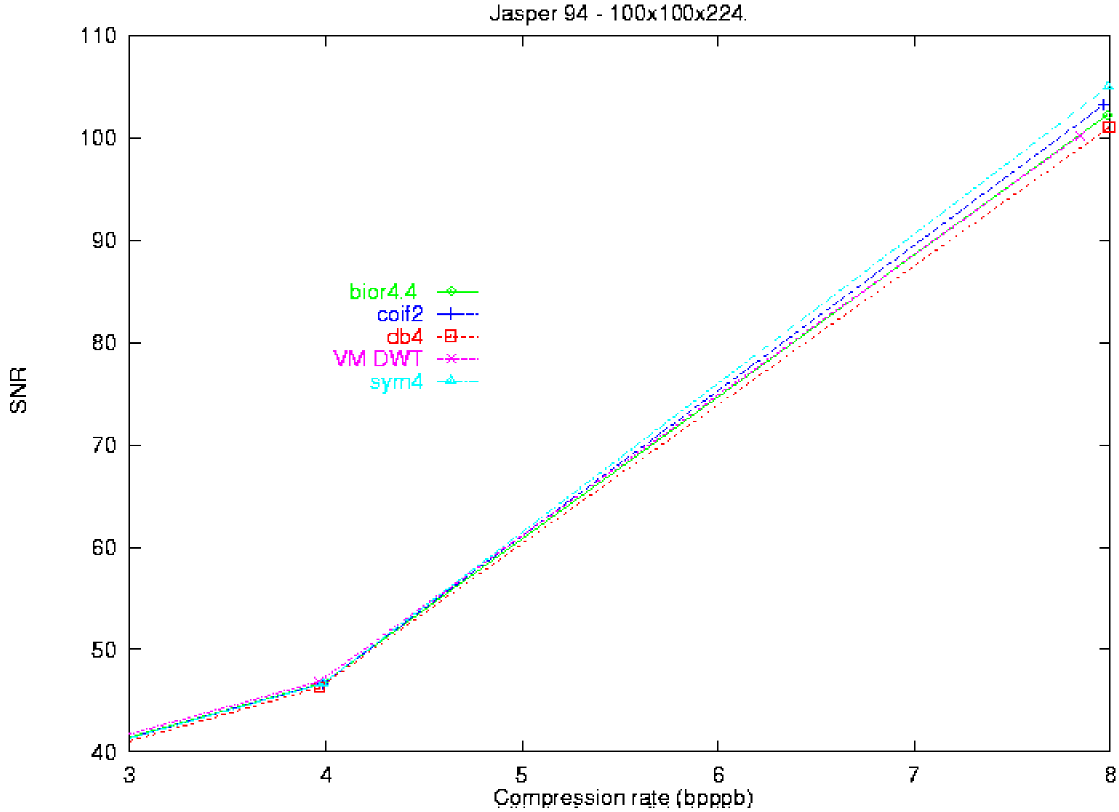

Figure 3. Jasper rate-distortion characteristics for rates above 2 bpppb.

3.2 AVIRIS decorrelation in the Fourier domain

This experiment uses the Jasper data set to emulate coding of Fourier transform spectroscopy (FTS) data. Such techniques are rapidly gaining prominence for both multi- and hyperspectral imaging because of the flexibility they offer for making finely-tuned spectral band selections, but unfortunately many of the currently operational systems are on military or intelligence platforms, making it difficult to obtain nonsensitive data sets for JPEG

experimentation. To get around this problem, we have simulated a FTS data set by discrete Fourier transforming the spectral pixel vectors in the Jasper data cube. Since it would be desirable, on a FTS platform, to transmit the raw data in its Fourier transform representation, we have performed our decorrelation and compression experiments on the Fourier transform representation of the Jasper cube.

A magnitude plot of a Jasper DFT vector is shown in Figure 4. The plot is symmetric since the input vector is real-valued. Such representations are known as “interferograms” in FTS jargon, referring to the way the FTS data is acquired by sensor systems, and we will refer to our DFT representations of Jasper data as “interferograms” even though they were generated by DFT’ing the Jasper data in the laboratory.

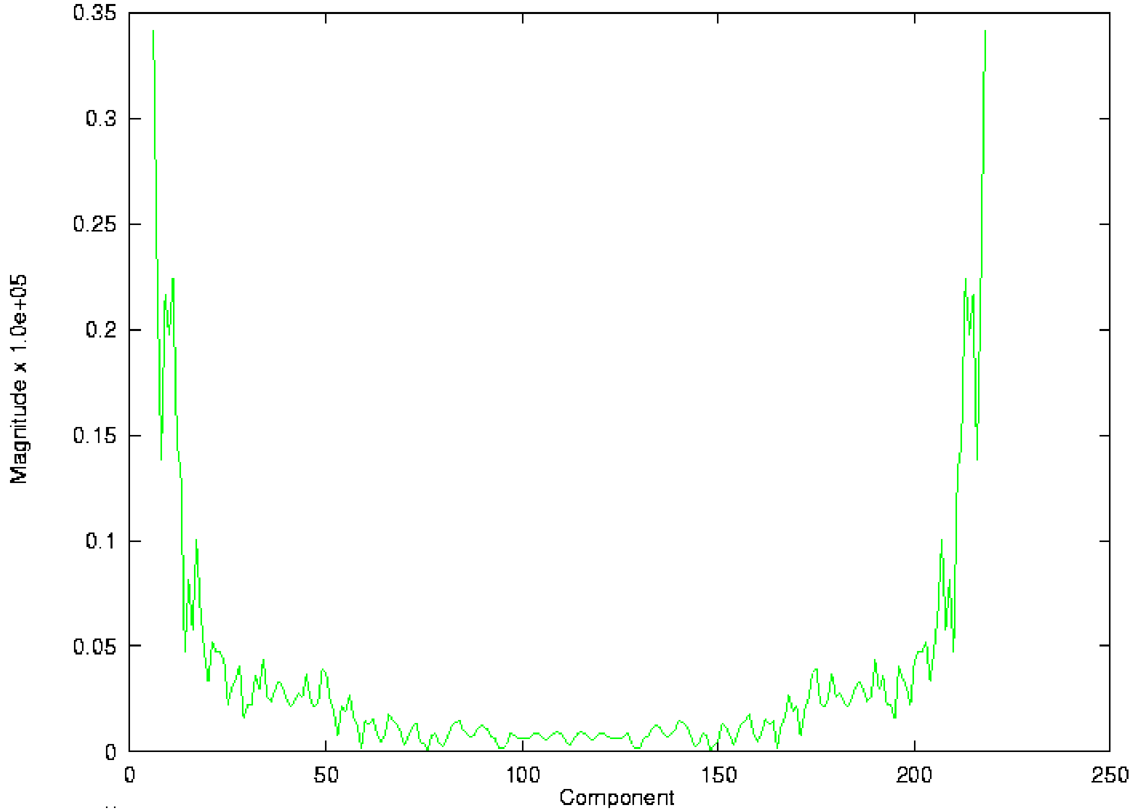

Figure 4. Magnitude plot of a Jasper interferogram.

As with the previous experiments, all floating point data generated outside the VM was quantized to 24 bits in order to input it into the VM for compression (and for component mixing in the case of the 9-7 filter bank implemented via symmetric extension). A block diagram is shown in Figure 5.

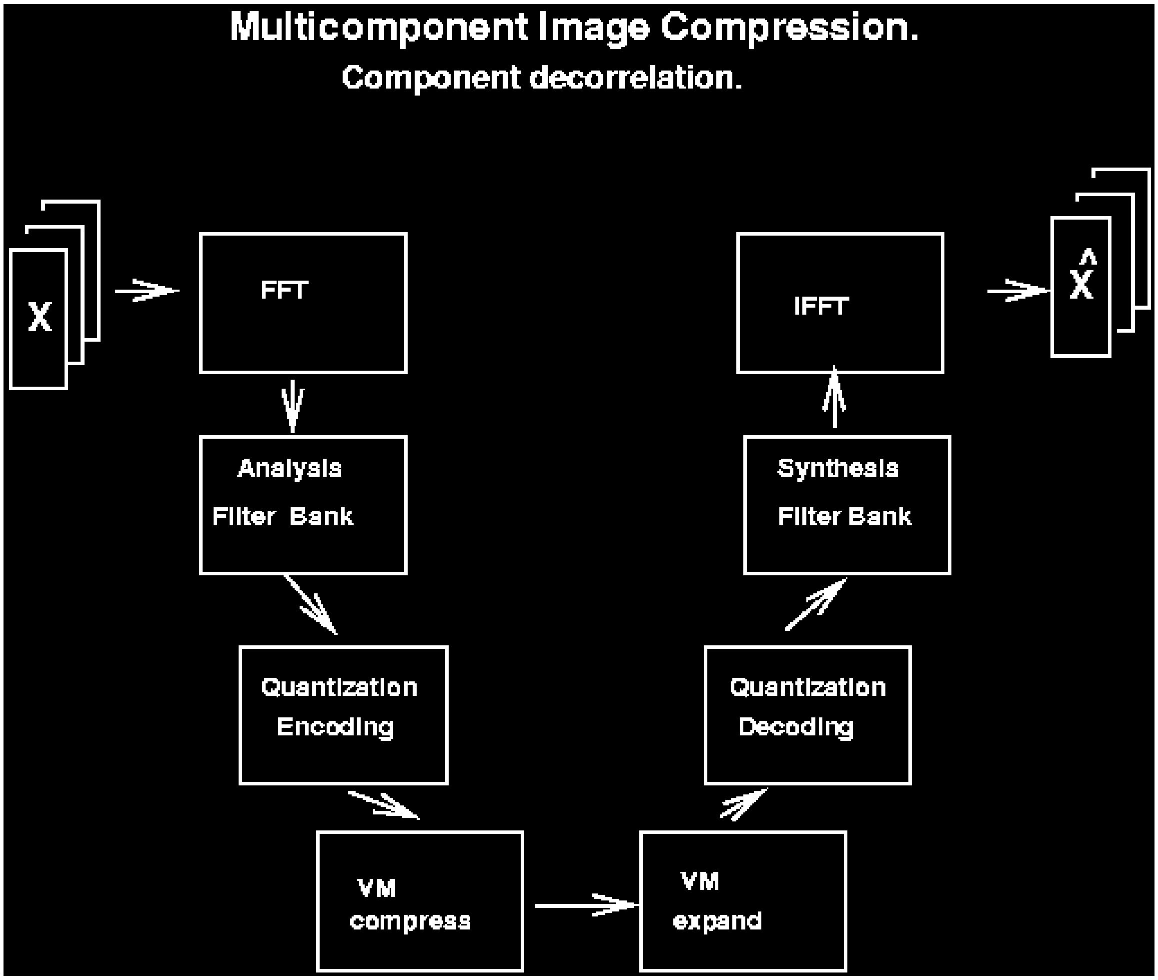

Figure 5. Block diagram for interferogram compression experiments.

Distortion measurements were made at 2 points in this diagram: in the Fourier domain and in the original Jasper image domain. Decorrelation and compression was performed on interferograms (i.e., in the Fourier domain) in both cases. Rate-distortion characteristics are plotted in Figure 6 and Figure 7. Relative performance is strikingly similar in both cases. The 9-7 implemented by symmetric extension is between half a dB and one dB better than the 9-7 implemented by circular convolution, and the performance of the symlet falls in between the two 9-7 curves in both the Fourier and image domains. In particular, the symlet is within a few tenths of a dB of the 9-7 symmetric and a few tenths better than the 9-7 circular, at all rates tested.

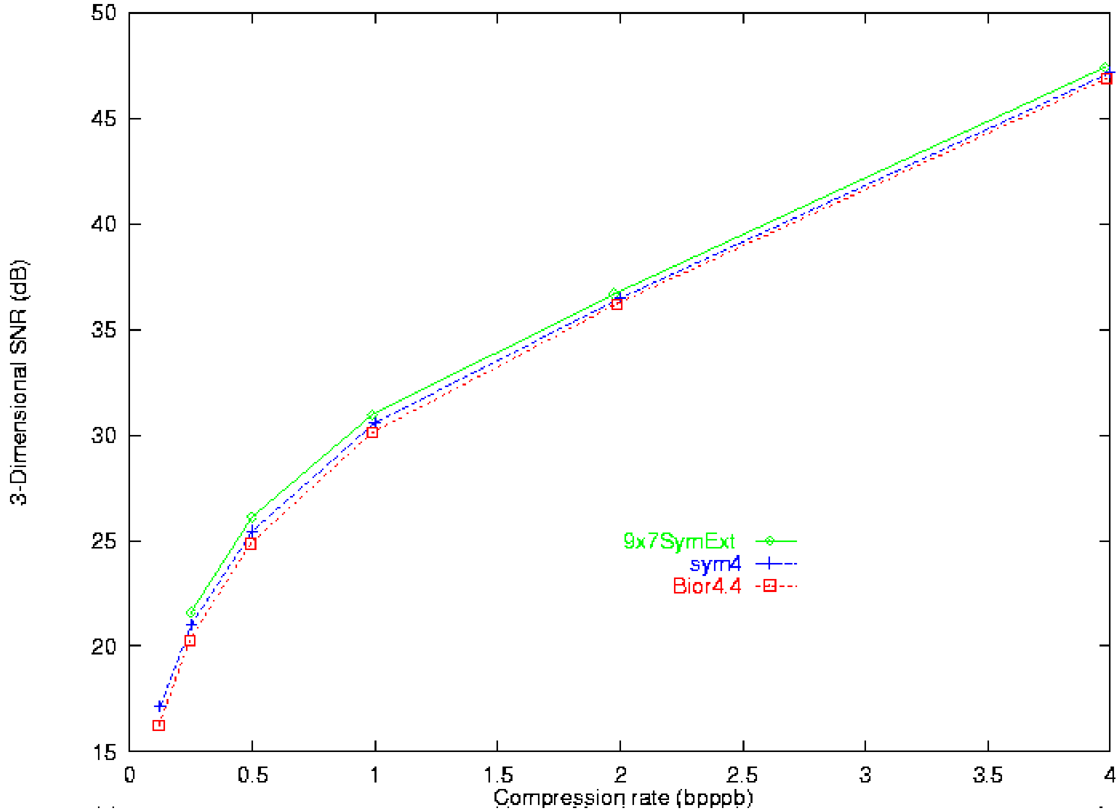

Figure 6. Fourier domain rate-distortion performance for interferogram compression.

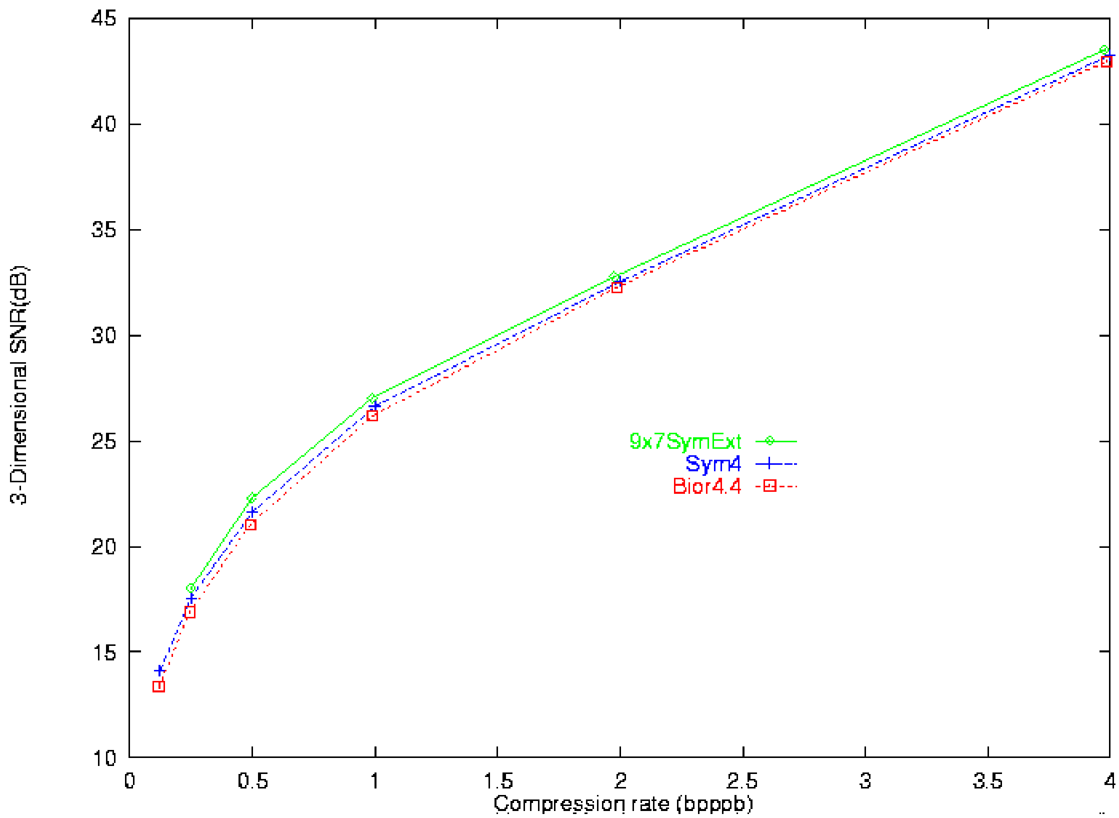

Figure 7. Image domain rate-distortion performance for interferogram compression.

4 Conclusions and Recommendations

The above results demonstrate that, for hyperspectral image component decorrelation, the 4-moment orthogonal symlet generates quantitative rate-distortion performance within experimental uncertainty of the performance of the 9-7 WS filter bank implemented with symmetric extension. Moreover, there are 2 significant factors that can be easily identified as biasing the results of these experiments against the symlet: the fact that it has only been implemented so far using circular convolution, and the necessity of quantizing the symlet-decorrelated components in order to pass them into the VM. Thus, we conclude that there is no quantitative basis for claiming that the 9-7 filter bank necessarily outperforms the symlet in hyperspectral component decorrelation applications. Moreover, orthogonal filter banks may be desirable to applications designers for reasons other than compression (e.g., feature extraction, image segmentation, reduced spectral resolution reconstruction, or pixel classification). This indicates that orthogonal filter banks are potentially viable options for multicomponent decorrelation. Consequently, we recommend that Part 2 syntax for signaling user-defined lifting steps be general enough to enable the signaling of nonlinear phase filter banks and that algorithms for implementing such filter banks for component decorrelation be developed in the context of Part 2 Annex I. We leave it up to WG1 to decide whether further core experiments are needed in this area.