

Compression and progressive transmission of astronomical images

Richard L. White

*Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218
rlw@stsci.edu*

Jeffrey W. Percival

*University of Wisconsin, Space Astronomy Laboratory, 1150 University Ave., Madison, WI 53706
jwp@sal.wisc.edu*

ABSTRACT

An image compression algorithm has been developed that is well-suited to astronomical images. The method has 3 steps: an intensity mapping to generate an image that has roughly constant noise in each pixel, an orthonormal wavelet transform, and quadtree coding of the bit-planes of the wavelet coefficients. The quadtree values may be further compressed by any standard compression technique, such as Huffman or arithmetic coding. If the 2-dimensional Haar transform is used, the calculations can be carried out using integer arithmetic, and the method can be used for both lossy and lossless compression. The Haar transform basis functions are well-suited to most astronomical images because they are highly localized. The performance of the algorithm using smoother, longer range wavelets is also shown; they can give slightly better lossy compression at the cost of an increase in artifacts around point sources, but they are not effective for lossless compression using this scheme.

This technique has also been used as the basis of a progressive image transmission system that can be used for either remote observing or access to remote image archives. After less than 1% of the data have been received, the image is visually similar to the original, so it is possible to assess the quality of images very quickly. If necessary, the entire compressed data set can be sent so that the original image is recovered exactly.

1. INTRODUCTION

Astronomical images consist largely of empty sky. Compression of such images can reduce the volume of data that it is necessary to store (an important consideration for large scale digital sky surveys) and can shorten the time required to transmit images (useful for remote observing or remote access to data archives.) Astronomical images can be extremely large, making the potential gains from image compression very important. For example, the Space Telescope Science Institute has digitized photographic plates covering the entire sky, generating 1500 images each having 14000×14000 16-bit pixels. Several astronomical groups are now constructing cameras with mosaics of large CCDs (each 2048×2048 or larger); these instruments will be used in projects that generate data at a rate exceeding 100 MBytes every 5 minutes for many years.

Astronomical images have some unusual characteristics that make most existing image compression techniques either ineffective or inapplicable. A typical image consists of a nearly flat background sprinkled with point sources and occasional extended sources. The images are often noisy, so that lossless compression does not work very well; furthermore, the images are usually subjected to stringent quantitative analysis, so any lossy compression method must be proven not to discard useful information, but must instead discard only the noise.

An example may make clear the difficulties of astronomical image compression. One of the simplest data compression techniques is run-length coding, in which runs of consecutive pixels having the same value are compressed by storing the pixel value and the repetition factor. This method is used in the standard compression scheme for facsimile transmissions. Unfortunately, it is quite ineffective for lossless

compression of astronomical images because even though the sky is *nearly* constant, the noise in the sky ensures that only very short runs of equal pixels occur. The obvious way to make run-length coding more effective is to force the sky to be exactly constant by setting all pixels below a threshold (chosen to be just above the sky) to the mean sky value. However, then one has lost any information about objects close to the detection limit. One has also lost information about local variations in the sky brightness, which severely limits the accuracy of photometry and astrometry on faint objects. Worse, there may be extended, low surface brightness objects that are not detectable in a single pixel but that are easily detected when the image is smoothed over a number of pixels; such faint structures are irretrievably lost when the image is thresholded to improve compression.

This paper describes an image compression algorithm that is well-suited to astronomical images. The method has 3 steps: (1) an intensity mapping to generate an image that has roughly constant noise in each pixel, (2) an orthonormal wavelet transform, and (3) quadtree coding of the bit-planes of the wavelet coefficients. The quadtree values may be further compressed by any standard compression technique, such as Huffman or arithmetic coding. This method is much better than techniques that keep only the wavelet coefficients with the largest amplitudes.

If the 2-D Haar transform is used as the wavelet transform, the calculations can be carried out using integer arithmetic, and the method can be used for both lossy and lossless compression. The Haar transform basis function are well-suited to most astronomical images because they are highly localized, and it is possible to adjust the coefficients during decompression to reduce the blockiness that comes from using such functions. The performance of the algorithm using smoother, longer range wavelets is also shown; they can give slightly better lossy compression, but they are not effective for lossless compression using this scheme.

This method is being used by the Space Telescope Science Institute to compress digitized versions of the Palomar and ESO Sky Survey plates for distribution on CD-ROM. Images compressed to about 1.5 bits/pixel are equivalent to the original images under both visual inspection and quantitative analysis. Images compressed to 0.2 bits/pixel are still useful, though some of the faintest objects are lost at such high compression factors.

This technique has also been used as the basis of a progressive image transmission system that can be used for either remote observing or access to remote image archives. After less than 1% of the data have been received, the image is visually similar to the original, so it is possible to assess the quality of images very quickly. If necessary, the entire compressed data set can be sent so that the original image is recovered exactly. It is also possible to speed the transmission even further by transmitting first only enough information to construct a version of the image that has been binned in blocks of 2×2 pixels; this is a natural feature of wavelet-based schemes.

2. THE H-TRANSFORM

The 2-dimensional Haar transform¹ (also known as the H-transform or the S-transform) can be used as the basis of an effective compression method for astronomical images²⁻⁵. The H-transform is calculated for an image of size $2^N \times 2^N$ as follows:

- Divide the image up into blocks of 2×2 pixels. Call the 4 pixels in a block a_{00} , a_{10} , a_{01} , and a_{11} .
- For each block compute 4 coefficients:

$$\begin{aligned}
 h_0 &= (a_{11} + a_{10} + a_{01} + a_{00})/2 \\
 h_x &= (a_{11} + a_{10} - a_{01} - a_{00})/2 \\
 h_y &= (a_{11} - a_{10} + a_{01} - a_{00})/2 \\
 h_c &= (a_{11} - a_{10} - a_{01} + a_{00})/2
 \end{aligned}
 \tag{1}$$

- Construct a $2^{N-1} \times 2^{N-1}$ image from the h_0 values for each 2×2 block. Divide that image up into 2×2 blocks and repeat the above calculation. Repeat this process N times, reducing the image in size by a factor of 2 at each step, until only one h_0 value remains.

This calculation can be easily inverted to recover the original image from its transform. The transform is exactly reversible using integer arithmetic if one is careful with the low-order bits of the coefficients⁵. It is straightforward to extend the definition of the transform so that it can be computed for non-square images that do not have sides that are powers of 2; the most effective way to do this is to assume reflected boundary conditions at the edges of the image. The H-transform can be performed in place in memory and is very fast to compute, requiring about $16M^2/3$ (integer) additions for a $M \times M$ image.

The H-transform can be derived from the 1-dimensional Haar transform, which involves taking sums and differences of pairs of adjacent elements in a vector. Apply a single sum/difference step of the 1-D transform along the rows of the images, then along the columns of the transformed image. Repeat this row/column transform, using only the sum coefficients (1/4 of the original image) as input. Repeat until only a single element remains.

2.1. Other wavelet transforms

The H-transform is a simple 2-dimensional discrete wavelet transform. The compression scheme described here is easily adapted for use with other wavelet transforms. Any 1-dimensional discrete wavelet transform can be converted to a 2-D transform as outlined in the last section, and the coefficients of that 2-D transform can be efficiently coded using the schemes described below. In this paper, compression results are also shown for an algorithm based on the Daubechies D4 wavelet transform⁶. The D4 transform has been modified so that it uses reflected boundary conditions rather than periodic boundary conditions.

The major advantage of the H-transform over the Daubechies and similar wavelet transforms is that the H-transform can be performed entirely with integer arithmetic, making it exactly reversible. Consequently it can be used for either lossless or lossy compression (as indicated below) and one does not need a special technique for the case of lossless compression (as was required, e.g., , for the JPEG compression standard and by FITSPRESS⁷.) However, the smoothness afforded by higher-order transforms can be advantageous.

3. QUANTIZATION

If the image is nearly noiseless, the H-transform is somewhat easier to compress than the original image because the differences of adjacent pixels (as computed in the H-transform) tend to be smaller than the original pixel values for smooth images. Consequently fewer bits are required to store the values of the H-transform coefficients than are required for the original image. For very smooth images the pixel values may be constant over large regions, leading to transform coefficients that are zero over large areas.

Noisy images still do not compress well when transformed, though. Suppose there is noise σ in each pixel of the original image. Then from propagation of errors, the noise in each of the H-transform coefficients is also σ . To compress noisy images, divide each coefficient by $S\sigma$, where $S \sim 1$ is chosen according to how much loss is acceptable. This reduces the noise in the transform to $0.5/S$ (because the largest error is 1/2 the least significant bit of the quotient), so that large portions of the transform are zero (or nearly zero) and the transform is highly compressible.

Why is this better than simply quantizing the original image? As discussed above, if we divide the image by σ then we lose all information on objects that are within 0.5σ of sky in a *single* pixel, but that are detectable by averaging a *block* of pixels. On the other hand, in dividing the H-transform by σ , we preserve the information on any object that is detectable by summing a block of pixels! The quantized H-transform preserves the mean of the image for every block of pixels having a mean significantly different than that of neighboring blocks of pixels.

If the noise is not constant across the image then this quantization method must be modified. The best approach we have found is to first scale the data to force the noise to be approximately constant in each pixel, and then to apply the H-transform and quantization described above. For CCD data, for example, the noise is a combination of Poisson counting statistics and readout noise. If we replace the input image I_{ij} by a scaled image $U_{ij} = 2\sqrt{I_{ij} + N^2}$, where N is the readout noise in each pixel, then the image U_{ij} has noise $\sigma_U \simeq 1$ in each pixel. U can then be compressed efficiently using the method described in this paper. Unfortunately, the use of this method for lossless compression is rather messy because considerable effort is required to make the square root transformation exactly reversible.




Figure 1. 256×256 pixel section of image from digitized Palomar Sky Survey E-plate of Coma galaxy cluster. Image has 16-bits; a logarithmic gray scale is used to make the noise near the sky brightness level visible.

4. QUADTREE CODING

The quantized H-transform has a rather peculiar structure. Not only are large areas of the transform image zero, but the non-zero values are strongly concentrated in the lower-order coefficients. The best approach we have found to code the coefficient values efficiently is quadtree coding of each bit-plane of the transform array. Quadtree coding has been used for many purposes⁸; the particular form we are using was suggested by Huang and Bijaoui⁹ for image compression.

- Divide the bit-plane up into 4 quadrants. For each quadrant code a ‘1’ if there are any 1-bits in the quadrant, else code a ‘0’.
- Subdivide each quadrant that is not all zero into 4 more pieces and code them similarly. Continue until one is down to the level of individual pixels.

This coding (which Huang and Bijaoui call “hierarchical 4-bit one” coding) is obviously very well suited to the H-transform image because successively lower orders of the H-transform coefficients are located in successively divided quadrants of the image.

We follow the quadtree coding with a fixed Huffman coding that uses 3 bits for quadtree values that are common (*e.g.*, 0001, 0010, 0100, and 1000) and uses 4 or 5 bits for less common values. This reduces the final compressed file size by about 10% at little computational cost. Slightly better compression can be achieved by following quadtree coding with arithmetic coding¹⁰, but the CPU costs of arithmetic coding are not, in our application, justified for 3–4% better compression. We have also tried using arithmetic coding directly on the H-transform, with various contexts of neighboring pixels, but find it to be both computationally inefficient and not significantly better than quadtree coding.

For completely random bit-planes, quadtree coding can actually use more storage than simply writing the bit-plane directly; in that case we just dump the bit-plane with no coding.




Figure 2. Effect of compression by H-transform and quadtree coding scheme described in this paper. This is also a sequence of images using the progressive transmission scheme; each image is the result of coding and transmitting another bit-plane from the H-transform.

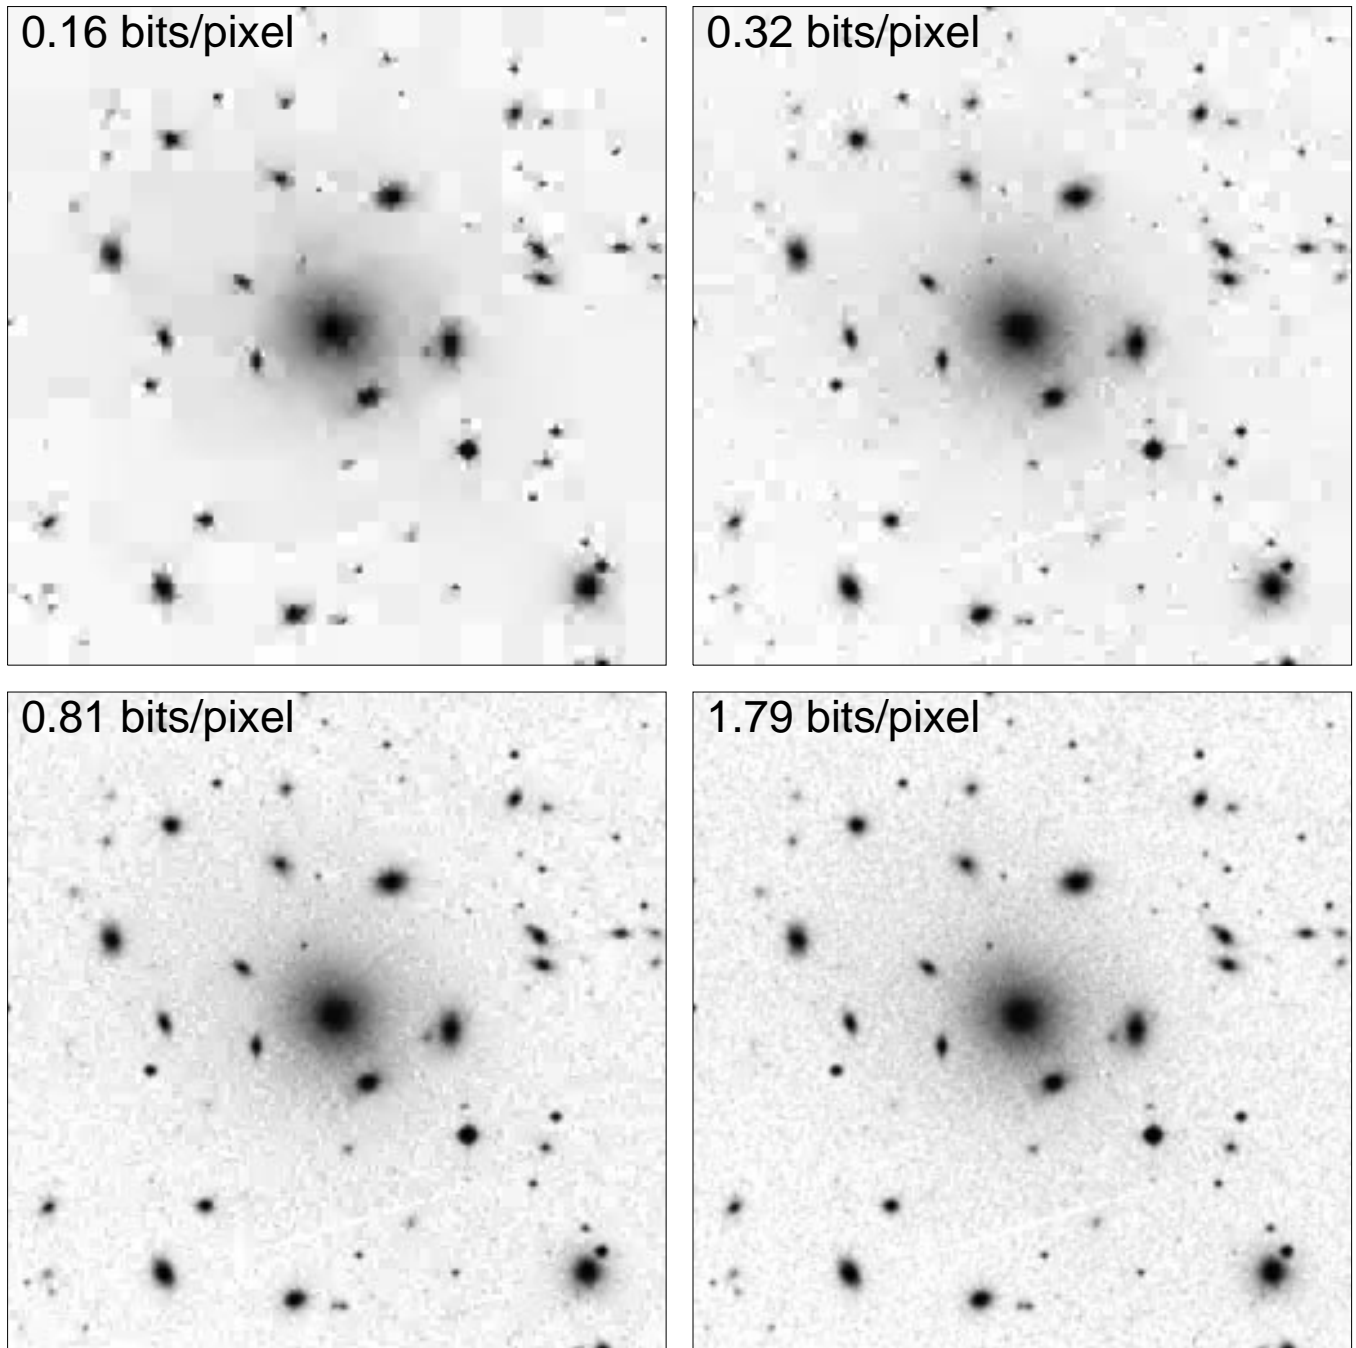


Figure 3. Images from Fig. 2 decompressed using adaptive smoothing method to reduce blocking artifacts.

5. EXAMPLES

As an example, Figure 1 shows a 256×256 section (7.2×7.2 arcmin) from a digitized version of the Palomar Observatory–National Geographic Society Sky Survey plate containing the Coma cluster of galaxies. This is a 16-bit image with noise $\sigma \simeq 315$ in each pixel. Figure 2 shows the resulting image for $S\sigma = 256, 512, 1024,$ and 2048 . These images are compressed by factors of 9, 20, 49, and 98 using the quadtree coding scheme. In all cases a logarithmic gray scale is used to show the maximum detail in the image near the sky background level; the noise is clearly visible in Figure 1. The image compressed by a factor of 9 is hardly distinguishable from the original. In quantizing the H-transform we have adaptively filtered the original image by discarding information on some scales and keeping information on other scales. This adaptive filtering is most apparent for high compression factors, where the sky has been smoothed over large areas while the images of stars have hardly been affected.

The adaptive filtering is, in itself, of considerable interest as an analytical tool for images⁴. For example, one can use the adaptive smoothing of the H-transform to smooth the sky without affecting objects detected above the (locally determined) sky; then an accurate sky value can be determined by reference to any nearby pixel.

The blockiness that is visible in Figure 2 is the result of difference coefficients being set to zero over large areas, so that blocks of pixels are replaced by their averages. It is possible to eliminate the blocks by an appropriate filtering of the image. A simple but effective filter can be derived by adjusting the H-transform coefficients as the transform is inverted to produce a smooth image; as long as changes in the coefficients are limited to $\pm S\sigma/2$, the resulting image will still be consistent with the quantized H-transform. The adaptively smoothed images corresponding to those in Figure 2 are shown in Figure 3.

5.1. Astrometric and photometric properties of compressed images

Astronomical images are not simply subjected to visual examination, but are also subjected to careful quantitative analysis. For example, for the image in Figure 1 one would typically like to do astrometric (positional) measurements of point sources to an accuracy much better than 1 pixel, photometric (brightness) measurements of objects to an accuracy limited only by the detector response and the noise, and accurate measurements of the surface brightness of extended sources.

We have done some experiments to study the degradation of astrometry and photometry on the compressed images compared to the original images¹¹. Even the most highly compressed images have very good photometric properties for both point sources and extended sources; indeed, photometry of extended objects can be improved by the adaptive filtering of the H-transform⁴. Astrometry is hardly affected by the compression for modest compression factors (up to about a factor of 20 for our digitized photographic plates), but does begin to degrade for the most highly compressed images.

These results are based on tests carried out with tools optimized for the original images; it is likely the best results will be obtained for highly compressed images only with analysis tools specifically adapted to the peculiar noise characteristics of the compressed images.

6. PROGRESSIVE TRANSMISSION

Note that by coding the transform one bit-plane at a time, the compressed data can be viewed as an incremental description of the image. One can initially transmit a crude representation of the image using only the small amount of data that is required for the sparsely populated, most significant bit-planes. Then the lower bit-planes can be added one by one until the desired accuracy is achieved. This could be useful, for example, if the data is to be retrieved from a remote database — one could examine the crude version of the image (retrieved very quickly) and abort the transmission of the rest of the data if the image is judged to be uninteresting. A version of this progressive transmission algorithm has been developed for the Wisconsin-Indiana-Yale-NOAO (WIYN) telescope¹². It will be used for remote observing and transfer of data over networks.

The improvement of the image during the progressive transmission is visible in Figure 2, where each successive panel represents the addition of another bit-plane from the H-transform. For comparison,




Figure 4. Progressive transmission of image by coding bit-planes of original image (rather than H-transform) using quadtree coding. Results are much inferior to those H-transform algorithm shown in Fig. 2.




Figure 5. Progressive transmission of image using Daubechies D4 discrete wavelet transform instead of H-transform. Results are slightly better than those shown in Fig. 2 for smooth objects, but show ringing artifacts around bright stars that might be objectionable for some applications.

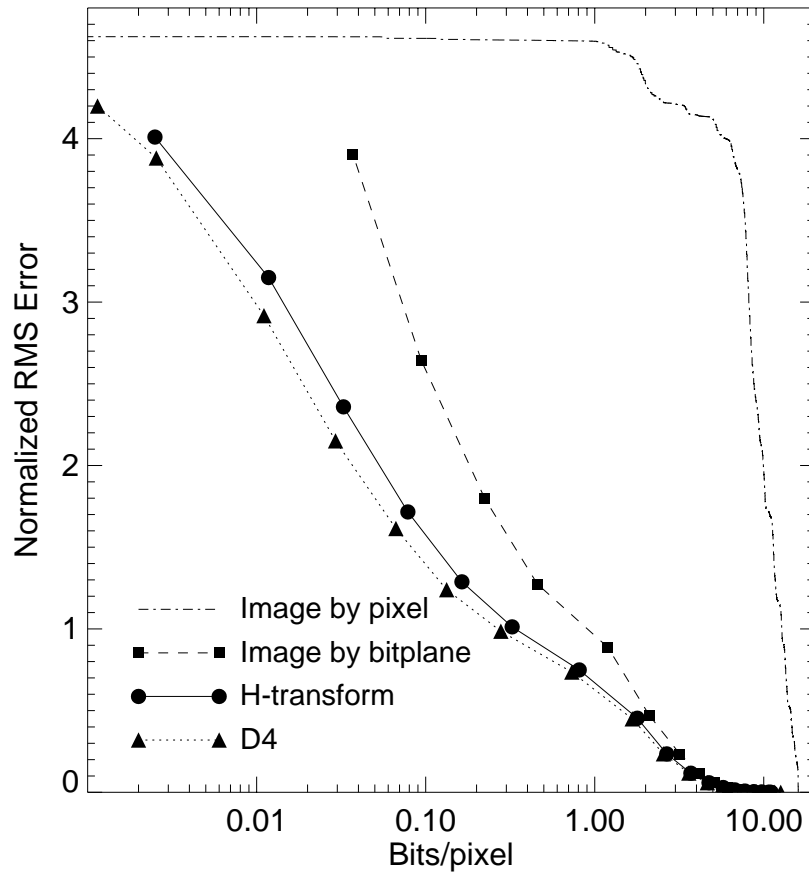


Figure 6. Mean square error in compressed images as a function of compression factor for various progressive transmission methods. The top curve shows the results from simply transmitting the pixels one by one with no compression. Other methods use quadtree coding on bit-planes of the image itself, the H-transform, and the D4 transform.

Figure 4 shows the performance of a progressive transmission scheme based on quadtree coding of the bit-planes of the image itself rather than a transformed version of the image. This method is a big improvement over simply sending the image pixel by pixel with no compression, but the results are not nearly as good as those obtained using the H-transform.

On the other hand, it may be possible to improve on the H-transform using other wavelet transforms. The results of progressive transmission using the Daubechies discrete wavelet transform D4⁶ are shown in Figure 5. Here the results appear slightly better than for the H-transform, especially for smooth objects. This is expected since the D4 wavelets form a smoother basis than the Haar functions. However, the D4-compressed images do show some ringing around bright stars (visible as “holes” or “ears” around the stars).

The D4 transform has been used previously for astronomical image compression by Press⁷, though in that application he simply retained the largest wavelet coefficients rather than coding bit-planes. (Press’s version of the two-dimensional D4 transform also has the drawback that it includes basis functions that are high frequency, hence localized, in one direction but low frequency, hence global, in the other direction. The D4 transform used here has only approximately isotropic basis functions.) Keeping the largest coefficients exactly is less efficient than keeping bit-planes of *all* coefficients; the difference is especially noticeable when one attempts to construct high fidelity compressed images, where the “largest coefficients” method requires

2–3 times as much data as the bit-plane method.

Figure 6 summarizes the performance of various algorithms using the mean square difference from the original image as a measure of image quality. This measure does not always correlate well with visual quality, but it is appropriate for these astronomical images where quantitative analysis is important. The figure shows the RMS error (normalized by the noise in the original image) versus the compression, expressed as the number of bits per pixel required to store the image. A normalized RMS error of unity means that the original and compressed images are consistent to about the noise level in the data.

Sending the image pixel-by-pixel is never competitive with these methods. Quadtree coding of the original image results in reasonably good results for lossless compression (11.12 bits/pixel) but poor quality for the early versions of the image constructed from the first few bit-planes. The D4 image is slightly better than the H-transform image at high compressions (in agreement with the visual assessment), but is considerably worse for lossless compression (12.60 bits/pixel for D4 versus 11.07 bits/pixel for H-transform).

Interestingly, the best lossless compression of these images is achieved by quadtree coding the difference of adjacent pixels along rows of the image (10.76 bits/pixel). However, these difference coefficients are absolutely useless for lossy compression because a slight change in any difference translates to a long streak across the image.

7. CONCLUSIONS

The algorithms described in this paper, based on either the H-transform or the Daubechies D4 wavelet transform, have been shown to be capable of producing highly compressed images that are very faithful to the original. Algorithms designed to work on the original images can give comparable results on object detection, astrometry, and photometry when applied to the images compressed by a factor of 10 or more. Further experiments will determine more precisely just what errors are introduced in the compressed data; it is possible that certain kinds of analysis will give more accurate results on the compressed data than on the original because of the adaptive filtering of the H-transform⁴.

The D4-compressed images are slightly superior to the H-transform images at high compression factors, though the D4 images do show some artifacts that might cause trouble for some image analysis programs. For lossless compression the H-transform method is better. For progressive transmission the H-transform leads to a very clean implementation that does not require any residual image to be transmitted to get a perfectly reconstructed image.

This work was supported by grant NAGW-2166 from the Science Operations Branch of NASA headquarters. The Space Telescope Science Institute is operated by AURA with funding from NASA and ESA.

REFERENCES

1. Haar, A. 1910, *Math. Ann.* **69**, 331.
2. Fritze, K., et al. 1977, *Astronomische Nachrichten*, **298**, 189.
3. Richter, G. M. 1978, *Astronomische Nachrichten*, **299**, 283.
4. Capaccioli, M., et al. 1988, *Astronomische Nachrichten*, **309**, 69.
5. Blume, H., and Fand, A. 1989, *SPIE Vol. 1091, Medical Imaging III: Image Capture and Display*, p. 2.
6. Daubechies, I. 1988, *Comm. Pure and Appl. Math.*, **41**, 909.
7. Press, W. H. 1992, in *Astronomical Data Analysis Software and Systems I*, eds. D. M. Worrall et al., (San Francisco: Astronomical Society of the Pacific), p. 3.
8. Samet, H. 1984, *ACM Computing Surveys*, **16**, 187.
9. Huang, L, and Bijaoui, A. 1991, *Experimental Astronomy*, **1**, 311.
10. Witten, I. H., Radford, M. N., and Cleary, J. G. 1987, *Communications of the ACM*, **30**, 520.
11. White, R. L., Postman, M., and Lattanzi, M. G. 1992, in *Digitized Optical Sky Surveys*, eds. H. T. MacGillivray & E. B. Thomson (Amsterdam: Kluwer), p. 167.
12. Percival, J. W., and White, R. L. 1993, in *Astronomical Data Analysis Software and Systems II*, eds. R. J. Hanisch et al., (San Francisco: Astronomical Society of the Pacific).