The Effects of Wavelet Compression on Digital Elevation Models (DEMs)

Michael J. Oimoen, Environmental Scientist, SAIC[1]
U.S. Geological Survey, EROS Data Center
47914 252nd St.
Sioux Falls SD 57198

Abstract **- This paper investigates the effects of lossy compression on floating-point digital elevation models using the discrete wavelet transform. The compression of elevation data poses a different set of problems and concerns than does the compression of images. Most notably, the usefulness of DEMs depends largely in the quality of their derivatives, such as slope and aspect.**

Three areas extracted from the U.S. Geological Survey's National Elevation Dataset were transformed to the wavelet domain using the third order filters of the Daubechies family (DAUB6), and were made sparse by setting 95 percent of the smallest wavelet coefficients to zero. The resulting raster is compressible to a corresponding degree. The effects of the nulled coefficients on the reconstructed DEM are noted as residuals in elevation, derived slope and aspect, and delineation of drainage basins and streamlines. A simple masking technique also is presented, that maintains the integrity and flatness of water bodies in the reconstructed DEM.

INTRODUCTION

We consider the lossy wavelet compression of floating-point digital elevation models, where the source data are areas selected from the U.S. Geological Survey's National Elevation Dataset (NED) [1]. The discrete wavelet transform has been used successfully in the field of image compression. Although image formats typically are confined to integer values, the same methods may be applied to floating-point data. DEMs differ substantially from image data in their application. While it is possible to view a DEM as an image by assigning a color or intensity to elevations, it is more common for a DEM to be used as a component in a broader analysis. Derived data, such as slope, aspect, and shaded relief are often as important, or more important, than the elevations themselves. Thus the compression losses will be examined both in terms of effects on absolute elevation and on elevation derivatives.

THE DISCRETE WAVELET TRANSFORM

The discrete wavelet transform (DWT) operates on vectors or matrices whose dimensions are powers of two, recursively decomposing them into high- and low-frequency components at increasingly coarser scales. The result is a matrix of the same size as the original, composed of nested submatrices [2].

By itself, the DWT is lossless and completely reversible, but an opportunity for significant compression is present because many of the values in the transformed matrix are small in amplitude, and consequently carry little information [3]. By setting a certain percentage of the smallest values to zero, the matrix is made sparse, and can be effectively compressed to a corresponding degree (e.g. a matrix with 80 percent of its coefficients set to zero can typically be compressed by 80 percent) using many common compression algorithms. This process is known as *hard thresholding* since the matrix values are either set to zero or left alone, and is the simplest form of wavelet compression [4]. Reversing the transform yields an approximation of the original data, often with remarkably high accuracy for the degree of compression that it allows. Figure 1 follows the progressive decomposition of a 512×512 pixel DEM through three iterations of a simple wavelet transform.

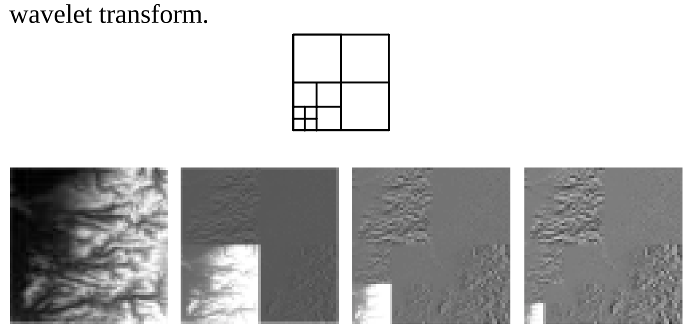

Figure 1. Raster DEM decomposed through three iterations of the Haar Transform

Effective compression also depends on the choice of *basis function.* Unlike Fourier transforms, which use sines and cosines exclusively, there are a wide variety of wavelet bases from which to choose. Figure 1 shows an implementation of the Haar transform, which is a pair of very simple low-pass and high-pass kernels, often used for illustrative purposes. More suitable for DEMs, and many other types of data are the Daubechies family of wavelets, which frequently are referred to as DAUB4, DAUB6, DAUB8, etc, where the numeric values correspond to the number of coefficients in their kernels. The selection of kernels involves balancing the retention of small features against smoothness in the reconstructed data [3].

[1] The work described in this article was performed under the USGS contract number 03CRCN0001.

DAUB6, which is a third order filter, was selected for this study. Nearly 95 percent of the wavelet coefficients must be set to zero before the effects of information loss are evident in a shaded relief image. Figure 2 depicts a full decomposition by DAUB6 and the locations of the coefficients that are retained after hard thresholding at 95 percent.

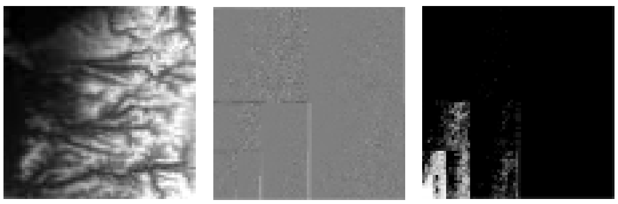

Figure 2. Raster DEM fully decomposed into DAUB6 wavelet space, and made sparse by setting 95% of largest coefficients to zero.

METHODOLOGY

Three areas of dissimilar terrain were selected from the NED, the locations of which are shown in figure 3. Each tile is composed of 4096×4096 pixel single-precision floating-point elevation values, and covers slightly over one square degree. Slope, aspect, and shaded relief were derived.

Figure 3. Locations of test data

The raw data were then transformed into 2-D wavelet space using DAUB6 filter coefficients. The transformed matrices were then made sparse by hard threshold criteria, setting 95 percent of the lowest magnitude values to zero. The inverse transform was applied to reconstruct an approximation of the original elevation data, and shaded relief, slope, and aspect were derived.

To evaluate the loss of information, the original and reconstructed elevation, slope, and aspect data were differenced (original – reconstructed), and the residuals were noted by statistical measure. Hypsometrically tinted shaded relief representations of the original and reconstructed data were created for qualitative visual comparison.

See appendix for detailed results.

ERROR COMPARED TO COMPRESSION

In comparing compression thresholds to the single metric of root-mean-square error, it is observed (Figure 4) that degradation occurs roughly exponentially, and that 95 percent compression is near the upper limit of what will produce acceptable results. The difference in magnitude between the curves has not been fully investigated, but depends on both the range of elevations and the complexity of terrain within each tile.

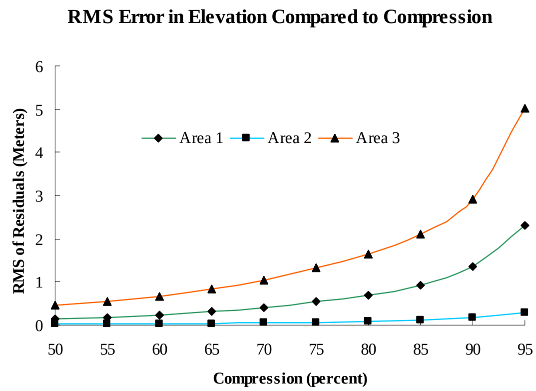

Figure 4. Graph of residual RMS compared to compression

HYDROLOGIC DERIVATIVES

The retention of hydrologic integrity is among the more critical demands that can be placed on any process that alters a DEM. Drainage basins and streamlines were derived for both the original and reconstructed versions of Area 3, and are submitted in Figure 5 for qualitative review.

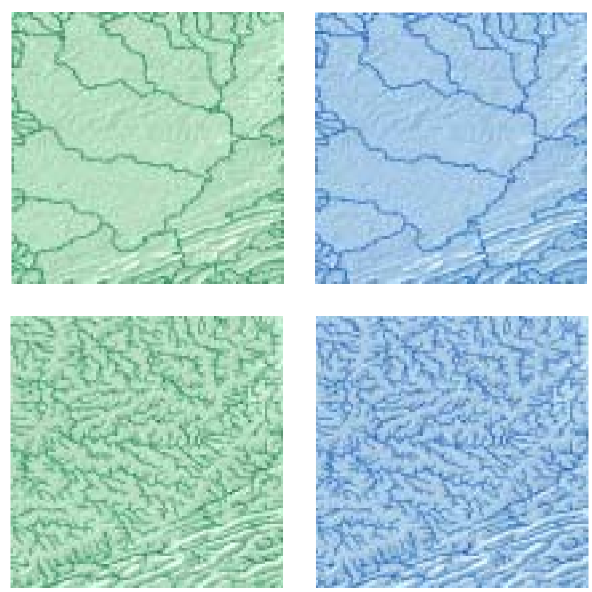

Figure 5. Watershed boundaries (top) and streamlines (bottom) derived from Area 3. Original data (left), and reconstructed from 5% of wavelet coefficients (right).

PRESERVATION OF FLAT AREAS

While the wavelet compression/reconstruction process has thus far been shown to be benign in it's effects on most derivatives, it can introduce spurious oscillations, or “ringing” artifacts into deliberately flattened areas of a DEM, such as water bodies. This is analogous to the Gibbs Effect encountered in the Fourier domain [5].

A simple but effective remedy, also developed independently in [5], is to extract flat areas from the original data prior to compression. Flat areas may be detected and delineated automatically by noting where derived aspect is undefined, or where local variance is zero. These criteria are used to create a binary mask that when intersected with the original DEM provides a raster of flat areas only, with their elevations intact (figure 6b).

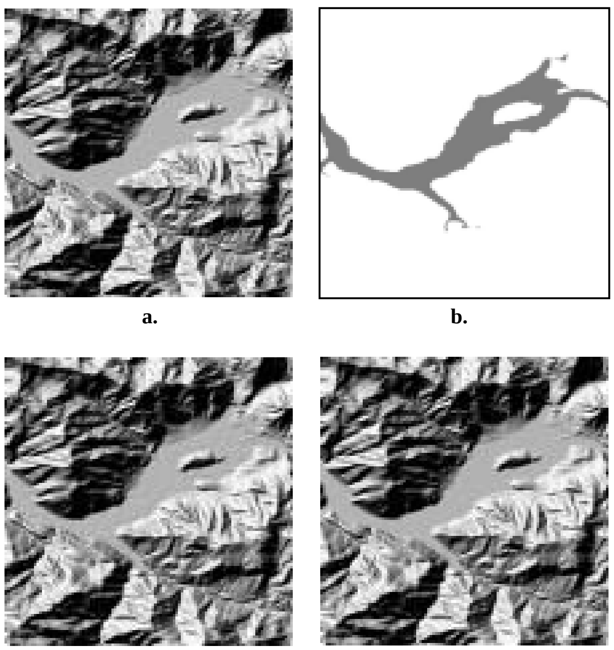

Figure 6. A priori masking of flat areas for reinsertion into reconstructed data. **a**. Detail of original data. **b**. Mask created from flat area. **c**. Detail of reconstructed data showing ringing in flat areas. **d**. Flat area restored by overlaying reconstructed data with flat mask.

This raster is highly compressible without any further transformation, as it is composed of a finite number of contiguous areas of constant elevation values. The 4096×4096 pixel mask created from Area 1 compresses to a mere 130 kilobytes, while retaining full floating-point accuracy. By retaining this data as a separate layer along with the compressed wavelet data, flat areas may be easily restored after the DEM is reconstructed (fig. 6d).

CONCLUSIONS

This paper demonstrates that the compression of floating point DEM data using the discrete wavelet transform produces largely acceptable results even in its simplest implementation, and can be enhanced to preserve flattened areas with minimal storage overhead. Other basis functions and compression techniques exist, and may produce even better results.

REFERENCES

[1] Gesch, D., M. Oimoen, S. Greenlee, C. Nelson, M. Steuck, D. Tyler, "The National Elevation Dataset", *Photogrammetric Engineering and Remote Sensing*, Volume 68, Number 1, January 2002, pp. 5-11.

[2] Jaffard, S., Y. Meyer, R. Ryan, *Wavelets: tools for science and technology*, SIAM (2001)

[3] Graps, A.L. "An Introduction to Wavelets", *IEEE Computational Sciences and Engineering*. Volume 2, Number 2, Summer 1995, pp 50-61.

[4] Press, W. H., S. Teukolsky, W. Vetterling, B. Flannery. *Numerical Recipes in C,* 2nd Edition, Cambridge University Press (1992)

[5] Bjørke, J.T., S. Nilsen "Wavelets applied to simplification of digital terrain models", International Journal of Geographical Information Science, Volume 17, Number 7, (2003), pp 601-621

APPENDIX

Area 1 is composed of a combination of flat and mountainous terrain, and is centered roughly on Green Peter Lake, Oregon.

Bounding Box		**Elevation Statistics**	
North	45.12°	Min	25.2 m
South	43.98°	Max	1809.4 m
East	121.98°	Mean	575.2 m
West	123.12°	StdDev	402.1 m

Area 2, in rural South Dakota near Huron, is largely flat, glacial terrain spotted with numerous small lakes.

Bounding Box		**Elevation Statistics**	
North	45.12°	Min	170.4 m
South	43.98°	Max	1433.9 m
East	97.98°	Mean	624.7 m
West	99.12°	StdDev	183.7 m

Area 3 samples the complex Appalachian terrain of West Virginia.

Bounding Box		**Elevation Statistics**	
North	38.12°	Min	368.8 m
South	36.98°	Max	650.5 m
East	80.98°	Mean	438.4 m
West	82.12°	StdDev	58.3 m

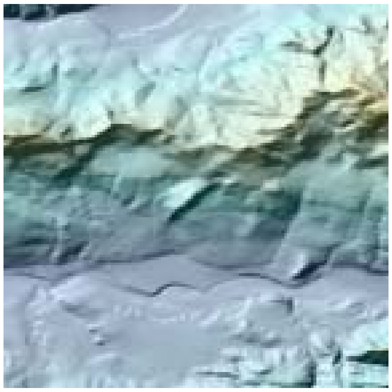

Detail, Area 1 (Original Data)

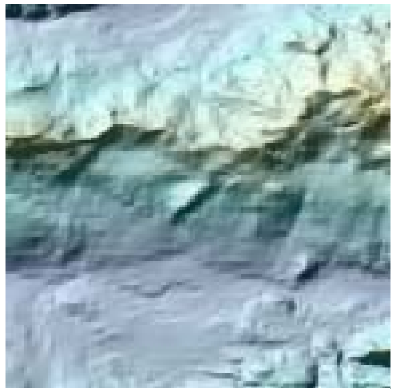

Reconstructed from 5% of wavelet coefficients.

Analysis of Residuals (Area 1)

Elevation		**Slope**	
Min	-20 m	Min	-17º
Max	20 m	Max	18º
Mean	0 m	Mean	0º
Std Dev	2.3 m	Std Dev	2.3º
Kurtosis	1.8	Kurtosis	1.6

Aspect Interquartile Range ±7º

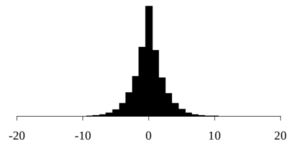

Elevation Residual (meters)

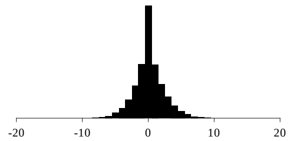

Slope residual (degrees)

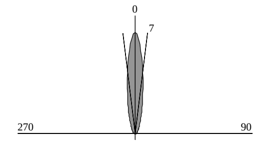

Aspect Residual (degrees)

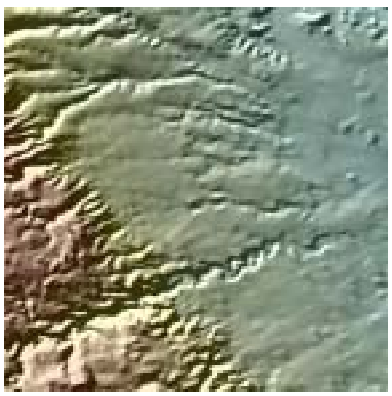

Detail, Area 2 (Original Data)

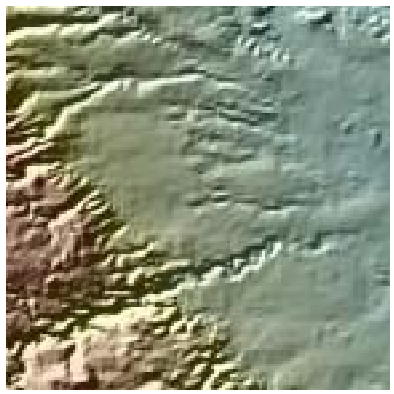

Reconstructed from 5% of wavelet coefficients.

Analysis of Residuals (Area 2)

Elevation		**Slope**	
Min	-3.1 m	Min	-2.4º
Max	2.8 m	Max	2.9º
Mean	0 m	Mean	0º
Std Dev	0.29 m	Std Dev	0.30º
Kurtosis	3.2	Kurtosis	3.3

Aspect Interquartile Range ±18º

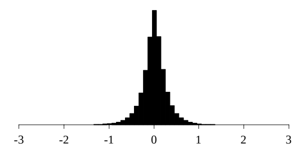

Elevation Residual (meters)

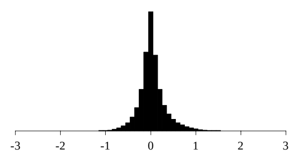

Slope residual (degrees)

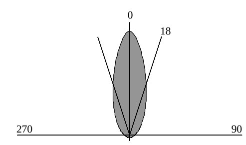

Aspect Residual (degrees)

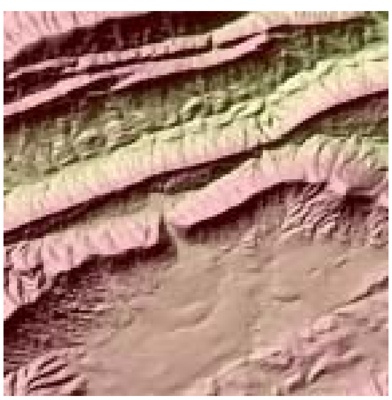

Detail, Area 3 (Original Data)

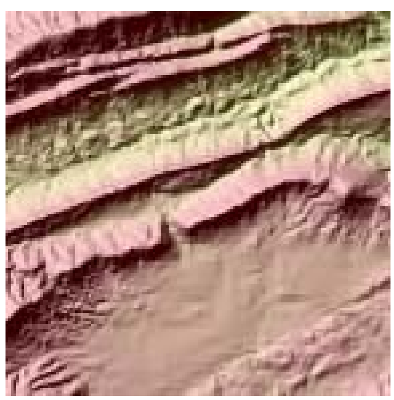

Reconstructed from 5% of wavelet coefficients.

Analysis of Residuals (Area 3)

Elevation		**Slope**	
Min	-33 m	Min	-26º
Max	34 m	Max	29º
Mean	0 m	Mean	0º
Std Dev	5.0 m	Std Dev	4.5º
Kurtosis	0.23	Kurtosis	0.42

Aspect Interquartile Range ±11º

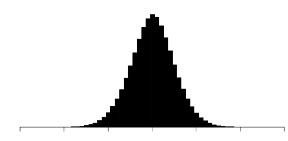

Elevation Residual (meters)

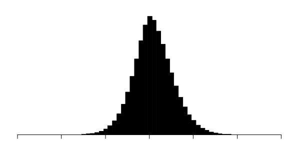

Slope residual (degrees)

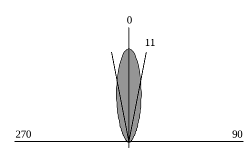

Aspect Residual (degrees)