ON THE USE OF JPEG 2000 TO ACHIEVE MINIMUM L-INFINITY ERROR WHEN SPECIFYING A COMPRESSION RATIO

Aldo Lucero[1], Sergio D. Cabrera[1], Edward Vidal Jr.[2]

[1]Dept. of Electrical and Comp. Eng., The University of Texas at El Paso, El Paso, TX 79968-0523
[2]Formerly at U.S. Army Research Lab., WSMR, NM 88002-5501

ABSTRACT

In this paper, minimization of the maximum absolute error (MAE) is achieved under a JPEG 2000 framework when a compression ratio is specified. The process is founded on our recent JPEG 2000-based algorithm for minimizing bit rate for a desired MAE using a *residual coding* approach [1-3] that uses the EBCOT coder separately. This type of algorithm uses the lossy and lossless capabilities of JPEG 2000 to achieve a desired MAE or a desired total bit rate. The technique to achieve the lowest MAE possible for a specified average bit rate is developed here for application to 3-D scientific data sets. Lossy compression is applied to the original data, while lossless compression is employed on the quantized residuals (the difference between the original and the lossy decompressed data). The lowest MAE is achieved by optimizing the allocation of the desired total bit rate between the two contributing rates corresponding to the lossy and the lossless compression steps. The methodology for achieving minimum MAE and results for 3-D meteorological data are presented here.

***Index Terms*—** JPEG 2000, 3-D coding, EBCOT, near-lossless, MAE, residual coding

1. INTRODUCTION

In scientific applications, data compression is usually employed in a lossless mode. The main reasons for avoiding lossy compression would be to stay away from any possible introduction of artificial artifacts or any possible loss of important data features. Nevertheless, in some applications, due to transmission or storage limitations, lossy compression can be used if the L_∞ error, also known as maximum absolute error (MAE), can be controlled. Furthermore, in some cases, the end user might only have the option of selecting the compression ratio or desired bit rate. For this scenario, achieving the lowest possible L_∞ error is desired. Here, a methodology to compress data using JPEG 2000 while achieving the lowest possible MAE for a specified compression ratio is presented. In [4] and [5], the lossy compression plus residual coding framework is also used. The data used in this study is meteorological data in GRIB format, which is a popular format used by the World Meteorological Organization (WMO). The data set used consists of 5 variables: DP (Dew Point), GPH (Geo Potential Height), T (Temperature), U and V (wind components). Each variable contains 20 slices, which have dimensions of 181 by 360 pixels each. Results obtained for this data set when specifying a desired MAE were presented previously in [3]. In this paper, the results shown correspond to specifying a desired bit rate while minimizing MAE.

2. DESCRIPTION OF COMPRESSION PROCESS

In this section, the compression process is discussed. The block diagram shown in Figure 1 summarizes the process. The first step is to lossy compress the original 3-D data using JPEG 2000 Part 2 as in [1-3, 7]. Our implementation of this approach is limited to the use of the Karhunen-Loeve transform (KLT) for decorrelation in the z-direction. We refer to the rate used to lossy compress the original data as the *open loop* rate (R_{OL}). The next step is to decompress in order to obtain the recovered data. Residuals, which are the difference between original and recovered data, are then computed next. Keeping and quantizing residuals is the key to being able to control the MAE or to minimize it for a specified bit rate. The quantizer used in this study is a zero-preserved symmetric uniform quantizer that uses $2^Q - 1$ levels. The quantizer step size Δ is determined from the largest residual magnitude and thus, by appropriately choosing Q, the final MAE is controlled. The number of bits needed to quantize the residuals has a direct impact on the bit rate needed to losslessly compress the quantized residuals. For the technique described here, the quantized residuals are compressed using the baseline (Part 1 of the standard) JPEG 2000 with zero levels of wavelet decomposition (i.e., no wavelet transform is used) and the *q-step reversible* option. This means that only the EBCOT [6] portion of JPEG 2000 is being used on the quantized residuals as a lossless coder. It should be mentioned that the original data has the mean subtracted from each slice. Moreover, the KLT transformation matrix must be saved in order to be able to perform the inverse in the decoding process. The vector of means and the transformation matrix comprise the overhead block shown in Figure 1. In the next section, the technique used to control the total bit rates while minimizing the MAE is discussed.

In order to take advantage of the energy compaction achieved from applying the KLT in the z-direction, an inter-slice bit-rate allocation technique is used. Following its successful presentation in [7], the JJ2K RDO (rate distortion optimization) bit-rate allocation technique is used in this study because it takes advantage of the rate-distortion information already generated inside the JJ2000 implementation of JPEG 2000. The RDO information is organized in an array where the index determines the lower bound of the slope in the MSE rate-distortion curve. This information is generated independently of the chosen target bit rate when a slice is compressed. Thus, by compressing the slice once, the complete derivative of the RDO curve can be formed. Furthermore, since the slopes vary monotonically with the index in the array of RDO information, the optimum bit-rate allocation can be easily calculated since the values of the slopes are discrete and fixed for all slices. Only the bit rates at which these values are achieved are different for each slice. An average bit rate for a candidate slope is calculated by simple averaging of the array values (bit rates) across the slices. Solving for the optimum allocation is done by computing all these averages and finding the correct one [7]. The only drawback is that the granularity of the possible RDO slope values produces a granularity in achievable average bit-rates. Examples of this downside are included in section 4.

3. CONTROLLING THE TOTAL BIT RATE

The methodology used to control the total bit rate while searching for the lowest possible MAE is described in this section. As discussed in section 2, the data goes through two compression steps. Thus, total bit rate is the sum of the open loop rate (R_{OL}) from the lossy compression of the original data plus the residual bit rate (R_{EBCOT}) needed to losslessly compress the quantized residuals. The constraint in searching for the lowest possible MAE is that $R_{OL}+R_{EBCOT}$ needs to be less than or equal to the desired total bit rate (DTR). The technique used in this study is based on an iterative exhaustive search, but as discussed below, two conditions are checked in order to reduce the number of iterations. The search technique is summarized as follows:

1. Initialize MAE_{Final} to a large value (e. g. 10^8)
2. Choose R_{OL} from values 0.05 to DTR in steps of 0.05.
3. For the current R_{OL}, perform the 3-D compression of the data and compute the resulting achieved open loop MAE (MAE_{OL}). Next, estimate the expected MAE from the quantizer by using the second loop, which corresponds to varying the number of bits (Q_{bits}) from 2 to 9. The expected MAE is computed as follows:
$$MAE_{Expected} = \frac{MAE_{OL}}{2^{Q_{bits}} - 2}$$
4. If the $MAE_{Expected}$ is less than MAE_{Final}, then compress/decompress the quantized residuals and set MAE_{Final} equal to the achieved MAE if R_{OL} plus R_{EBCOT} is less than DTR. Otherwise, there is no need to compress/decompress the residuals.
5. Also, if R_{OL} plus R_{EBCOT} is already greater than DTR, exit the second loop on Q_{bits}. This can be performed since R_{EBCOT} monotonically increases with respect to increasing Q_{bits} as shown in Figure 2.
6. Stop the R_{OL} loop when R_{OL} equals DTR or go to Step 2.

There will be cases where the sum of the two achieved rates will be much lower (~0.2 bpp) than DTR. This can be attributed to two factors. To a lesser degree, the granularity of the JJ2K RDO allocation causes the achieved open loop rates to be the same for several target R_{OL}'s. The biggest impact is due to the fact that R_{EBCOT} is limited to the few rates that can be achieved for each value of Q_{bits}. This problem is also a granularity issue since the achieved EBCOT rate can drastically increase just by going one bit higher. The results obtained using this algorithm are shown in the next section.

4. RESULTS

In this section, the results obtained for the algorithm just described are presented. Results for variables DP, T and U are summarized in Tables 1, 2 and 3 respectively. Similar results were obtained for variables GPH and V. For example, for variable GPH and a target total bit rate of 0.5, R_{OL} = 0.40 with R_{EBCOT} = 0.03 and a final MAE of 1.70 were obtained. At the same target total bit rate for variable V, R_{OL} = 0.10, R_{EBCOT} = 0.25, and a final MAE of 2.04 were achieved. For all 3 tables, it can be noted that the sum of the open loop rate and the EBCOT rate is always less than or equal to the target total. They also include a column where the number of bits used to quantize the residuals is shown.

Target Total	Open Loop Rate	EBCOT Rate	Achieved Total Rate	Bits	Final MAE
1.0	0.45	0.26	0.71	3	1.84
0.9	0.45	0.26	0.71	3	1.84
0.8	0.45	0.26	0.71	3	1.84
0.7	0.15	0.46	0.61	4	2.30
0.6	0.10	0.45	0.55	4	2.64
0.5	0.30	0.17	0.47	3	2.77
0.4	0.20	0.09	0.29	3	4.41
0.3	0.20	0.09	0.29	3	4.41
0.2	0.10	0.09	0.19	3	6.16
0.1	0.05	0.02	0.07	2	19.18

Table 1 – Results for DP

For the results shown in Table 1, it can be observed that the same solution was obtained for target bit rates from 1.0

to 0.8. For the cases where the answer yielded a total bit rate significantly lower than the target (1.0 and 0.9), further reduction of the final MAE can be achieved by taking advantage of the gap between target and achieved rates. For example, coding of *outliers* could be used to further reduce the final MAE (see Figure 3). The coding of these outliers could be as simple as the direct coding of their values and their positions.

For the results shown in Table 2, which correspond to the T variable, a similar analysis can be performed for target bit rate totals of 0.6 and 0.7. The same occurs for target total bit rates 0.4 and 0.5. Table 3 contains the results obtained for variable U. It can be noted again that the sum of R_{OL} and R_{EBCOT} is always lower than the target total. The repetition of the same answer as far as the total achieved rate can also be seen for this variable. As noted earlier, this leaves room for improvement where the difference between total achieved and total target can be used to further code more quantized residuals in order to decrease the final MAE.

Target Total	Open Loop Rate	EBCOT Rate	Achieved Total Rate	Bits	Final MAE
1.0	0.65	0.34	0.99	3	0.30
0.9	0.50	0.34	0.84	3	0.40
0.8	0.20	0.52	0.72	4	0.53
0.7	0.35	0.19	0.54	3	0.61
0.6	0.35	0.19	0.54	3	0.61
0.5	0.25	0.13	0.38	3	0.90
0.4	0.25	0.13	0.38	3	0.90
0.3	0.20	0.09	0.29	3	1.24
0.2	0.10	0.07	0.17	3	1.85
0.1	0.05	0.01	0.06	2	6.86

Table 2 – Results for T

Target Total	Open Loop Rate	EBCOT Rate	Achieved Total Rate	Bits	Final MAE
1.0	0.50	0.41	0.91	3	0.67
0.9	0.35	0.28	0.63	3	0.92
0.8	0.35	0.28	0.63	3	0.92
0.7	0.35	0.28	0.63	3	0.92
0.6	0.25	0.20	0.45	3	1.33
0.5	0.25	0.20	0.45	3	1.33
0.4	0.35	0.01	0.36	2	2.77
0.3	0.25	0.01	0.26	2	3.97
0.2	0.15	0.02	0.17	2	4.78
0.1	0.05	0.02	0.07	2	6.50

Table 3 – Results for U

The results shown in Table 4 have been included in order to verify in more detail the results for the DP variable. They were obtained by running the complementary algorithm that minimizes average bit-rate [1, 2] with a desired MAE bound set equal to 2.50. The way to verify the results of Table 1 is to observe that as the target total rate is decreased from 1.0 to 0.1, the last target rate to achieve a final MAE less than 2.50 corresponds to 0.7. As Table 1 indicates, the achieved total rate for this case corresponds to R_{OL}=0.15 and R_{EBCOT}=0.46 with a final MAE equal to 2.30. These values can be compared with those in Table 4 for the case where the total rate is the lowest. As expected, the lowest rate occurs when a target open loop rate of 0.15 is used. The needed EBCOT rate (to keep the MAE below 2.5) for this case is equal to 0.46 and the resulting MAE is equal to 2.30. These are the same results that were obtained in Table 1 for a total target rate of 0.7.

Another observation that can be made from the results in Table 4 is the effect of the granularity of the bit allocation technique used for the open loop compression. As mentioned before, the RDO slope information from the JPEG 2000 implementation is being extracted and used. The downside to this is that the resolution of the bit rates that may be selected is limited. A direct effect of this can be observed for target bit rates from 0.45 to 0.65. The actual rate achieved for all these cases was equal to 0.52. It should be remembered that the achieved rates correspond to the average bit rate from the rates assigned to the KLT slices during the bit allocation process.

Target Rates	Actual Rates	Final MAE	EBCOT Rate	Bits	Total Rate
0.05	0.06	1.28	1.07	5	1.12
0.10	0.08	1.23	1.06	5	1.15
0.15	0.14	2.30	0.46	4	0.60
0.20	0.21	1.89	0.54	4	0.75
0.25	0.21	1.89	0.54	4	0.75
0.30	0.35	1.19	0.83	4	1.18
0.35	0.35	1.19	0.83	4	1.18
0.40	0.35	1.19	0.83	4	1.18
0.45	0.52	1.84	0.26	3	0.78
0.50	0.52	1.84	0.26	3	0.78
0.55	0.52	1.84	0.26	3	0.78
0.60	0.52	1.84	0.26	3	0.78
0.65	0.52	1.84	0.26	3	0.78

Table 4 – Exhaustive Search for Desired MAE of 2.5

5. CONCLUSIONS

A technique for minimizing the maximum absolute error while keeping the total bit rate close to and bounded by some desired value DTR has been presented. This is done using JPEG 2000 as the only compression engine and is developed so that 3-D compression can be performed on scientific data, where L_∞ error is important. The results showed that this was achieved, and one example was verified by comparing it to the results obtained using the complementary algorithm that controls the desired MAE [1, 2]. Future work includes coding the outliers for the cases where there is room to get close to the total target rate. This includes looking at different techniques that could be used to code them. Finally, an algorithm based on modeling distortion versus bit rate is currently being developed in order to avoid the type of exhaustive search used here.

6. REFERENCES

[1] A. Lucero, S. D. Cabrera, and Ed Vidal Jr., "Residual Coding with a JPEG2000 Framework for Controlling L-Infinity Error Applied to Ultraspectral Sounder Data," SPIE Optics and Photonics 2006, Vol. 6300, San Diego, CA, 2006.

[2] A. Lucero, S. D. Cabrera, A. Aguirre, and Ed Vidal Jr., "Evaluating Residual Coding with JPEG2000 for L-Infinity Driven Hyperspectral Image Compression," SPIE Optics and Photonics 2005, Vol. 5889, San Diego, CA, 2005.

[3] Y. Yee, E. Vidal Jr., A. Lucero, A. Aguirre, and S. D. Cabrera, "Evaluating Optimal Techniques for Compressing NOGAPS Data Sets for Quick Response Applications," American Meteorological Society, 87th Annual Meeting, January 2007, San Antonio, TX.

[4] N. Wang and R. Brummer, "Experiment of a Wavelet-Based Data Compression Technique with Precision Control," 19th International Conference on Interactive Information Processing Systems (IIPS), Long Beach, CA, American Meteorological Society, 2003.

[5] S. Yea and A. Pearlman, "A Wavelet-Based Two-Stage Near-Lossless Coder," IEEE Transactions on Image Processing, Vol. 15, No. 11, pp. 3488-3500, November 2006.

[6] D. Taubman, "High Performance Scalable Image Compression with EBCOT," IEEE Transactions on Image Processing, Vol. 9, no. 7, pp. 1158-1170, July 2000.

[7] A. Aguirre, Ed Vidal Jr., and S. D. Cabrera, "Assessment of inter-slice bit-allocation strategies in the application of JPEG 2000 to ultraspectral sounder data," SPIE Optics and Photonics 2006, Vol. 6300, San Diego, CA, 2006.

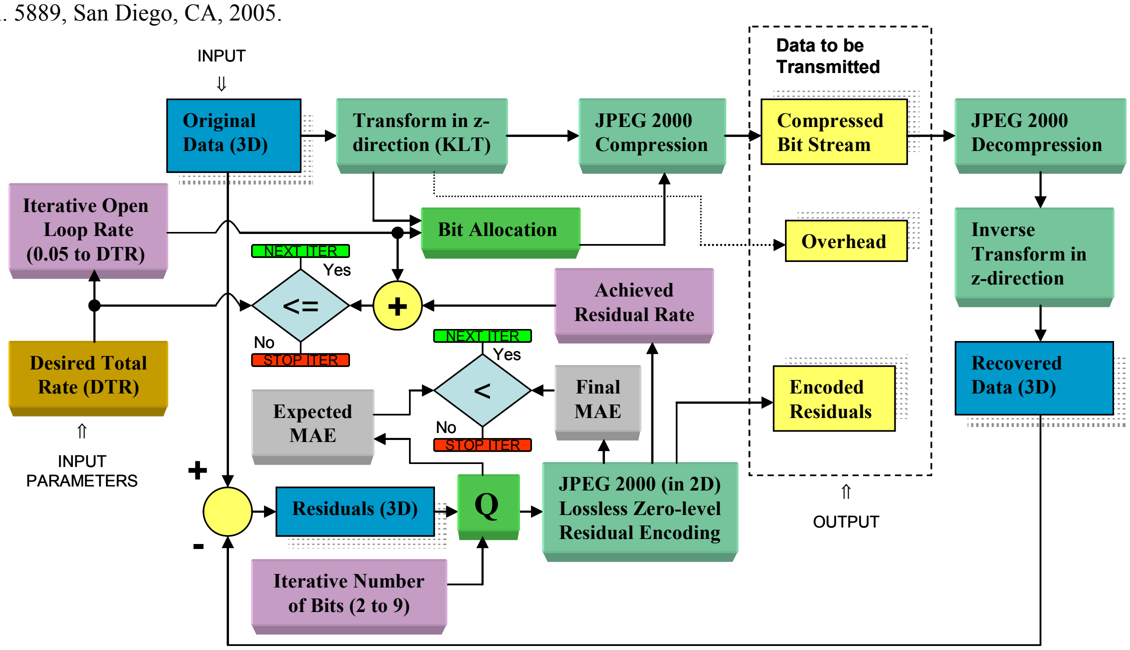

Figure 1 – Block diagram of compression process

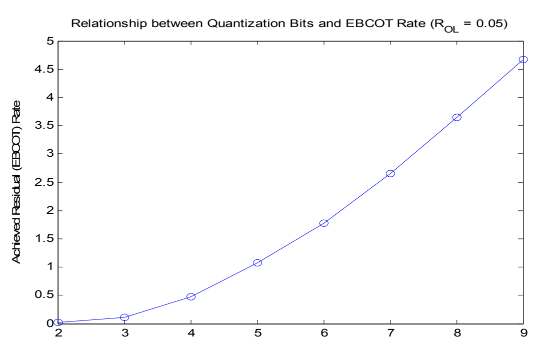

Figure 2 – Relationship between Q_{bits} and EBCOT Rate

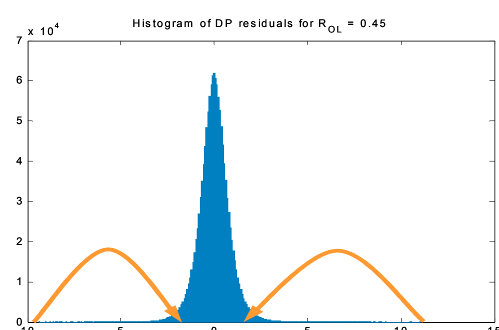

Figure 3 – The Concept of Encoding Outliers