Compression of 3-Dimensional Medical Image Data Using Part 2 of JPEG 2000

Alexis Tzannes, Ph.D. ‡
Aware, Inc.

Nov. 24, 2003

1. Introduction

JPEG 2000 is the new ISO standard for image compression commonly used to compress medical images. Part 1 of the standard provides the core coding system, specifying both lossy and lossless compression. Part 2 provides extensions to the standard useful for a variety of applications. For 3 dimensional data sets there are Part 2 extensions that allow for the use of several different types of decorrelating transformations in the third dimension, referred to as multiple component transformations (MCT). Specifically, wavelet transforms, linear transforms, and dependency transforms (such as DPCM techniques) are all specified under Part 2.

These multi-component transformations in Part 2 of JPEG 2000 can be effective in compressing volumetric datasets because the correlation between adjacent images can be exploited to achieve better compression than if each image were compressed independently. For lossless compression, the reversible 5-3 wavelet transform is applied in the third dimension. For lossy compression, the non-reversible, floating point 9-7 wavelet transform is applied.

This document addresses the compression of 3-dimensional, volumetric image data using Part 2 of the JPEG 2000 standard. It reports on the results of using the wavelet-based multi-component transforms defined in Part 2 to compress volumetric medical datasets. Test results for both lossless and lossy compression are presented.

---

‡ Alexis Tzannes is a senior member of the JPEG 2000 development team at Aware, Inc. He has been an active participant in the US and International JPEG 2000 standard body activities for the last 4 years. Within the international committee (ISO/IEC JTC1/ SC29/ WG1), he is the chair of the Ad Hoc Group **JPEG 2000 for Medical Imaging,** which supports and promotes the adoption of the JPEG 2000 standards in the medical imaging community.

2. Lossless Compression

For lossless compression of 3-dimensional datasets, the reversible 5-3 wavelet transform is first applied independently to each pixel in the third dimension. The 5-3 transform is an integer transform, so the data remains in integer format throughout this process. After applying this transform to each pixel, each image in the transformed sequence is compressed losslessly using Part 1 of JPEG 2000.

The results are compared to those achieved by applying Part 1 lossless compression to each image independently. Results are summarized in Section 5.1.

3. Lossy Compression

For lossy compression of 3-dimensional datasets, the irreversible 9-7 wavelet transform is first applied independently to each pixel in the third dimension. The 9-7 transform is a floating point transform, so the data is converted to floating point during this operation. After applying this transform to each pixel, each image of the series is compressed independently using the lossy option of Part 1 of JPEG 2000.

The compression results are compared to those achieved by applying Part 1 lossy compression to each image independently. Results are summarized in Section 5.2.

4. Component Collections

The multi-component transformations in Part 2 also allow for the grouping of components into independent collections (tiling in the third dimension). For each pixel, the multi-component transform is applied to the pixel samples in each component collection independently. The size of the component collection and the specific components that go into each collection are indicated in the header of the compressed file.

Using component collections may be advantageous for 3D imagery that is highly correlated over a small number of components, but less correlated over the entire sequence. Smaller component collections can also be used to limit the amount of memory needed to perform a complete 3D compression (or decompression) on a large data set. Finally, for the decoder, component collections can be used for improved access to the compressed imagery, as only a single collection must be decoded to decompress a single image in the sequence.

In Section 5 below, we present results of lossless and lossy compression, using component collections of 20, 40, 80 and N frames, where N is the total number of components in the given sequence.

5. Results

The 3D compression techniques described above were tested on three different sets of volumetric medical data. Characteristics of the image sequences are given in the table below:

Table 1. Test sequence information.

Sequence name	Image size	# of images (N)	Bitdepth	Uncompressed sequence size
Seq. #1	256 x 256	127	8	7.9 Mbytes
Seq. #2	512 x 512	449	16	224 Mbytes
Seq. #3	512 x 512	620	16	310 Mbytes

5.1 Lossless Compression Results

Each of the three sequences was compressed losslessly using Part 1 JPEG 2000 and Part 2 JPEG 2000 using components collections of 20, 40, 80 and N, where N is the total number of images in the sequence. The resulting compressed image sizes are shown in the table below.

Table 2. Lossless Compression Results. Note that Part 2 JPEG 2000 compression reduces the compressed sequence size by 15-18%.

Sequence name	Uncompressed sequence size (Mbytes)	2D Part 1 JPEG 2000 compressed size (Mbytes)	3D Part 2 JPEG 2000 compressed size, at different component collection sizes (Mbytes)				Improvement with 3D compression, collection size N
			20	40	80	N	%
Seq. #1	7.9	3.81	3.27	3.25	3.24	3.23	15.2
Seq. #2	224	75.8	62.7	62.3	62.1	61.8	18.5
Seq. #3	310	120	105	104	101	100	16.7

Reviewing the lossless results presented in the Table 2, the compressed file size is significantly smaller for the 3D case compared to simple Part 1 JPEG 2000 compression. The best compression is achieved when all the components in the sequence are grouped into a single collection (the N column in the table), although the loss in compression efficiency (compared to the N column) is small even with a component collection as small as 20. In the right most column is the percent reduction in file size between Part 1 JPEG 2000 compression and Part 2 JPEG 2000 with a single component collection (of all the images in the sequence). Overall, Part 2 reduces the compressed sequence size by 15-18%, for these three sequences.

5.2 Lossy Compression Results

Each of the three sequences was compressed with loss using both Part 1 JPEG 2000 and Part 2 JPEG 2000 for 2D-only and 3D compression, respectively. Compression ratios in the range of 10:1 to 50:1 were used, and the resulting compressed sequences were compared using an average peak signal to noise ratio (pSNR) metric over all the images in the sequence. pSNR is a commonly used metric for image fidelity; higher pSNR is equivalent to lower distortion in the image.

To test the effect of the component collection size on compressed sequence size and resulting image pSNR, the first sequence was compressed using components collections of 20, 40, 80, and 127 (all the images in the sequence) images. The resulting average pSNRs over the entire image sequence are shown in Fig. 1, at compression ratios ranging from 10 to 50 to 1. The different Part 2 curves represent different component collections. As in the lossless case, the best results are obtained with a single component collection for the entire sequence. The average pSNR is slightly lower for the other component collections, but still considerably higher than that for Part 1 2D-only compression.

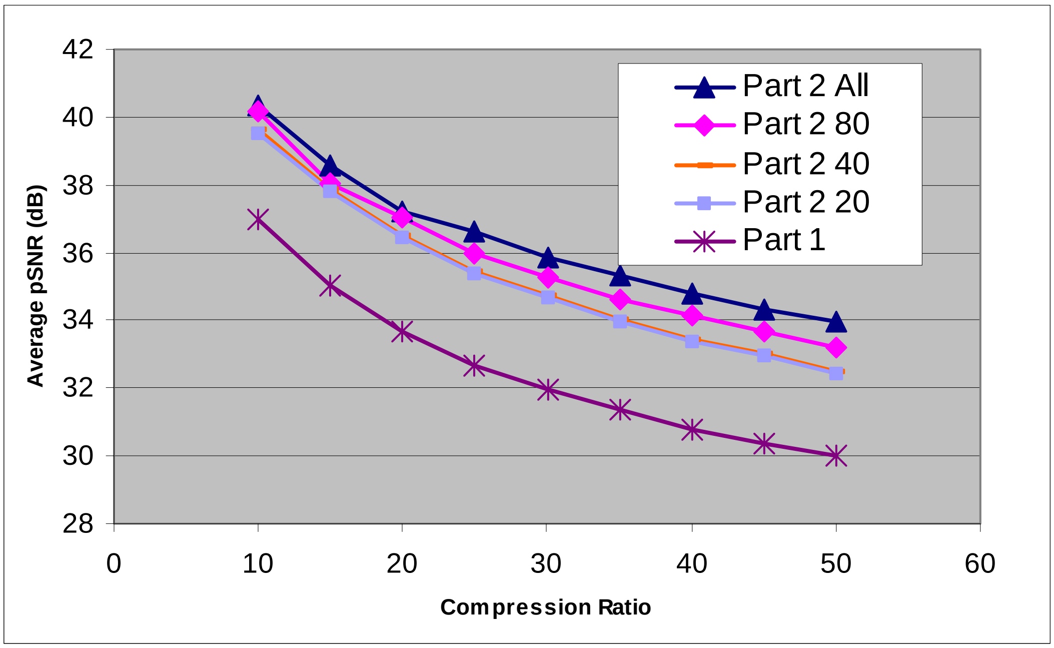

Figure 1. Sequence #1: average pSNR over all images in the sequence vs. compression ratio. The different Part 2 curves represent different component collections.

Results of compressing Sequences 2 and 3 are shown in Figures 2 and 3. These figures show the average pSNR of the resulting image sequence at various compression ratios, using a single component collection for the entire sequence. Note that the average pSNR is consistently and considerably higher than the Part 1 compression case.

Also note that at a specific pSNR level, Part 2 achieves more than twice the compression of Part 1. For example, for Sequence #2 in Figure 2, Part 1 compression results in an average pSNR of 77 dB, at 10:1 compression ratio. The same average pSNR can be achieved with Part 2 at a compression ratio of 25:1. Similarly, for Sequence 3, Part 1 achieves 67 dB at 15 to 1 compression ratio. Part 2 achieves the same pSNR value at 40 to 1 compression ratio. At higher compression ratios the gains are even larger, as seen in all three figures.

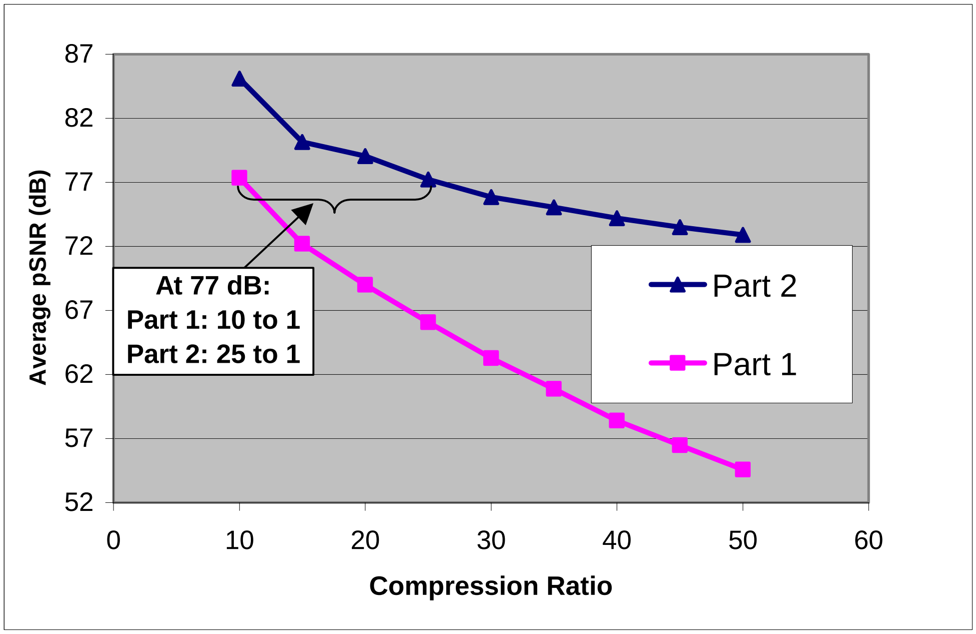

Figure 2. Sequence #2: Average pSNR over all images in the sequence vs. compression ratio. Note that Part 2 JPEG 2000 outperforms simple Part 1 JPEG 2000 by 8 to 18 dB. The Part 2 data in this figure was compressed using a single component collection for the entire sequence (429 images).

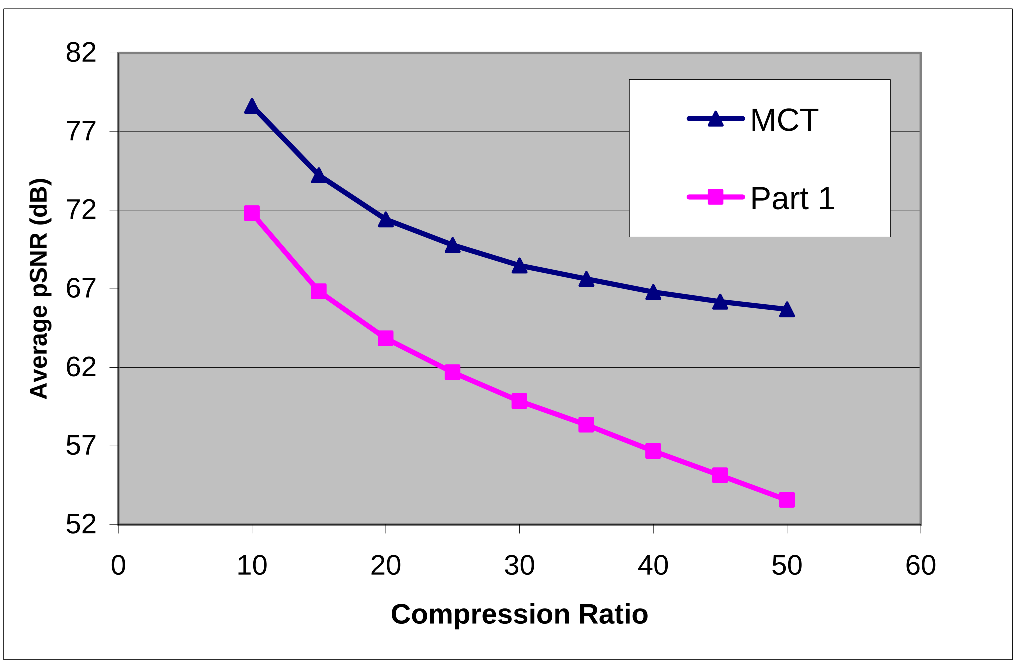

Figure 3. Sequence #3: Average pSNR over all images in the sequence vs. compression ratio. Note that Part 2 JPEG 2000 outperforms simple Part 1 JPEG 2000 by 6-13 dB. The MCT data in this figure was compressed using a single component collection for the entire sequence (620 images).

Figure 4 shows the pSNR of each image in Sequence 2, compressed at a compression ratio of 15 to 1. Notice that Part 2 displays consistent pSNR values around 80 dB, while Part 1 achieves pSNR values of about 72 dB. Similar results can be generated at other compression ratios, as indicated by the average pSNR values in Figures 1-3.

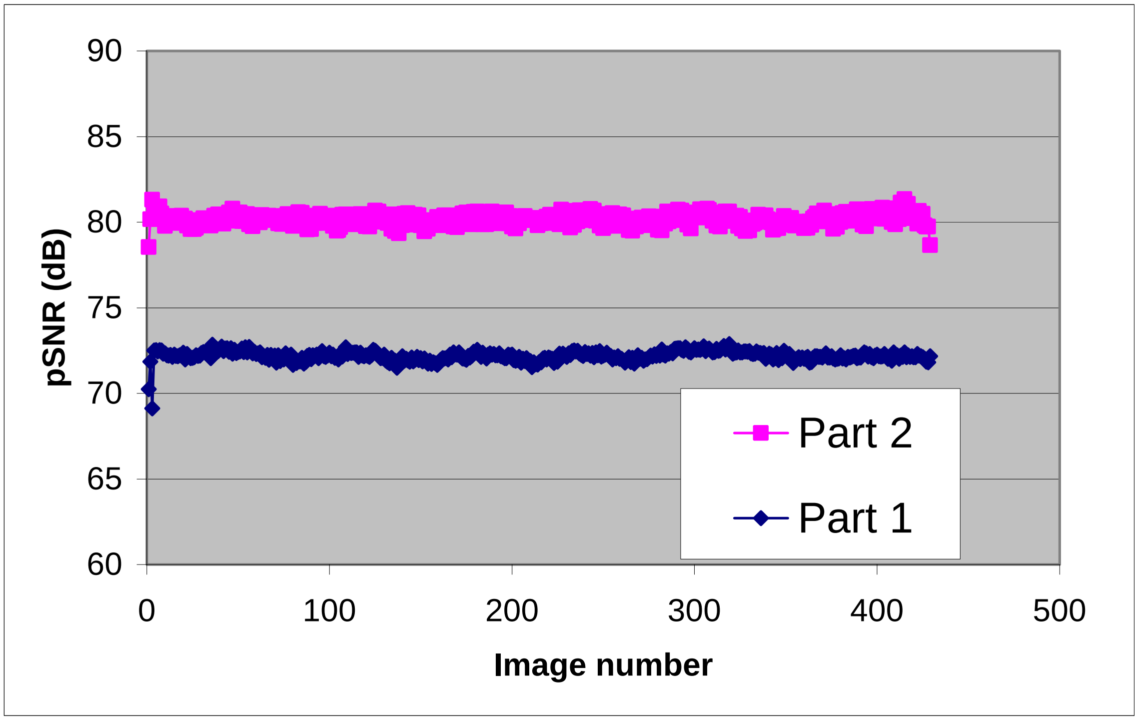

Figure 4. Sequence #2: Comparison of the pSNR of each image in the sequence at a compression ratio of 15 to 1. Note that Part 2 JPEG 2000 outperforms simple Part 1 JPEG 2000 consistently by about 8 dB. The Part data in this figure was compressed using a single component collection for the entire sequence (429 images).

6. Conclusions

This document reports on the performance of 3D data compression using Part 2 JPEG 2000 and Part 1 JPEG 2000. For lossless compression, applying Part 2 to take advantage of inter-image correlation achieved a compression ratio that was 15-18% higher than if images were compressed independently. For lossy compression, the results were even more dramatic. At a specific compression ratio, using Part 2 increased the average pSNR of the image sequence by 5-18 dB over the Part 1 case. This translates to an increase in compression of a factor of 2-3 for a given average pSNR value.