

Wavelet based ECW image compression

CHRISTOPH UEFFING, Merzhausen

ABSTRACT

The wavelet based ECW image compression is compared with older compression techniques and other wavelet compression methods. The ability to compress images without intermediate tiling or intermediate disk storage is a big advantage of the ECW compression especially for the compression of big remote sensing data sets.

The technique of ECW compression will be explained in more detail and typical applications for the compression are mentioned.

In a practical example some results of an ECW compression with different compression ratios of an airphoto are shown.

1. INTRODUCTION

In the last years the amount of data of remote sensing data sets grew up very fast. Not only the upcoming usage of digital airphotos but also the development of high resolution satellite systems (for example IKONOS) asked for a new approach of intelligent data storage and handling. Here the wavelet based ECW image compression can help to solve this problems.

2. COMPARING DIFFERENT COMPRESSION TECHNIQUES

2.1. Wavelet versus older compression techniques

Over the years, considerable effort has been spent in compressing information in order to make it easy to distribute, and to reduce storage requirements.

Techniques used by common compression program such as ZIP, or by older compression formats such as JPEG or TIFF, suffer from sever disadvantages when trying to compress large digital images. These problems include:

- Limited compression rates. Because ZIP and related methods are lossless techniques (it is important that a compressed program is exactly decompressed), they don't take advantage of the lossy compression techniques that result in much higher compression rates for digital imagery. Typically a "ZIPped" image file will be 50% smaller than the original image, whereas a wavelet compressed image will be 95% smaller than the original image.
- All or nothing decompression. It is essential to be able to selectively decompress a portion of digital imagery while viewing the image, and to selectively decompress the image at different levels of detail (as the user zooms in or out of the image view). Older techniques used by ZIP, TIFF, JPEG and other image compression formats, were not designed with selective decompression of imagery in mind. This means that they can not effectively be used when working with images larger than about half computer memory size, which is significantly smaller than the typical size of digital imagery today.
- Artefacts and visible errors in the data. Older techniques, for example as currently used by the JPEG compression process, compress the image as a series of blocks. This is because older techniques were memory based, so they needed to limit the size of each block. Because of this, JPEG and related formats suffer from significant and visible degradations when higher compression ratios are used. Recent breakthroughs in wavelet processing (such as the patent

pending method used by ER Mapper 6.0's Compression Wizard) have removed memory limits from image compression. This means that large imagery can be compressed efficiently, without introducing visible artefacts into the compressed image.

- No geographic coordinates. Older image formats do not have, or have very limited, geographic information stored with the compressed image. This is because older formats were designed for graphics art imagery rather than earth related digital imagery such as airphotos or satellite images. Because of this, many formats don't support geographic information, or only do so in a very limited way.
- Slow speed. In order to work effectively with large digital images, the user needs to be able to view any sub-section of the image, at any zoom factor, with sub-second response times. Older formats do not achieve adequate response rates for several reasons:
 1. Many formats require decompressing of the entire image, rather than a selected subset at a select level of detail.
 2. Many formats assume the image will be smaller than computer RAM, making them very slow to access large images.
 3. Most formats do not employ "clustering" of imagery information to reduce disk seeks.

2.2. Wavelet based image compression and decompression

Wavelet compression involves a way analyzing an uncompressed image in a recursive fashion, resulting in a series of higher resolution images, each "adding to" the information content in lower resolution images.

The primary steps in wavelet compression are performing a Discrete Wavelet Transformation (DWT), quantization of the wavelet-space image subbands, and then encoding these sub bands.

Wavelet images by and of themselves are not compressed images, rather it is the quantization and encoding stages that do the image compression. Image decompression, or reconstruction, is achieved by carrying out the above steps in reverse and inverse order. Thus, to restore the original image, the compressed image is decoded, dequantized, and then an inverse-DWT is performed.

Because wavelet compression inherently results in a set of multi-resolution images, it is well suited to working with large imagery which needs to be selectively viewed at different resolutions, as only the levels containing the required level of detail need to be decompressed.

Wavelet mathematics embraces an entire range of methods each offering different properties and advantages. For example, it is possible to compress 3 or more dimensional imagery using wavelets.

Wavelet compression has not been widely used because the DWT operation takes a lot of compute power, and because historical techniques perform the DWT operation in memory or by storing intermediate results on hard disk. This limits either the size of the image that can be compressed, or the speed at which it can be compressed.

The ER Mapper 6.0 ECW wavelet compression wizard uses a breakthrough new technique for performing the DWT and inverse-DWT operations (patent pending), which makes the use of wavelet based compression a practical reality.

The following diagram shows different applications for ER Mapper Compressed Wavelet imagery:

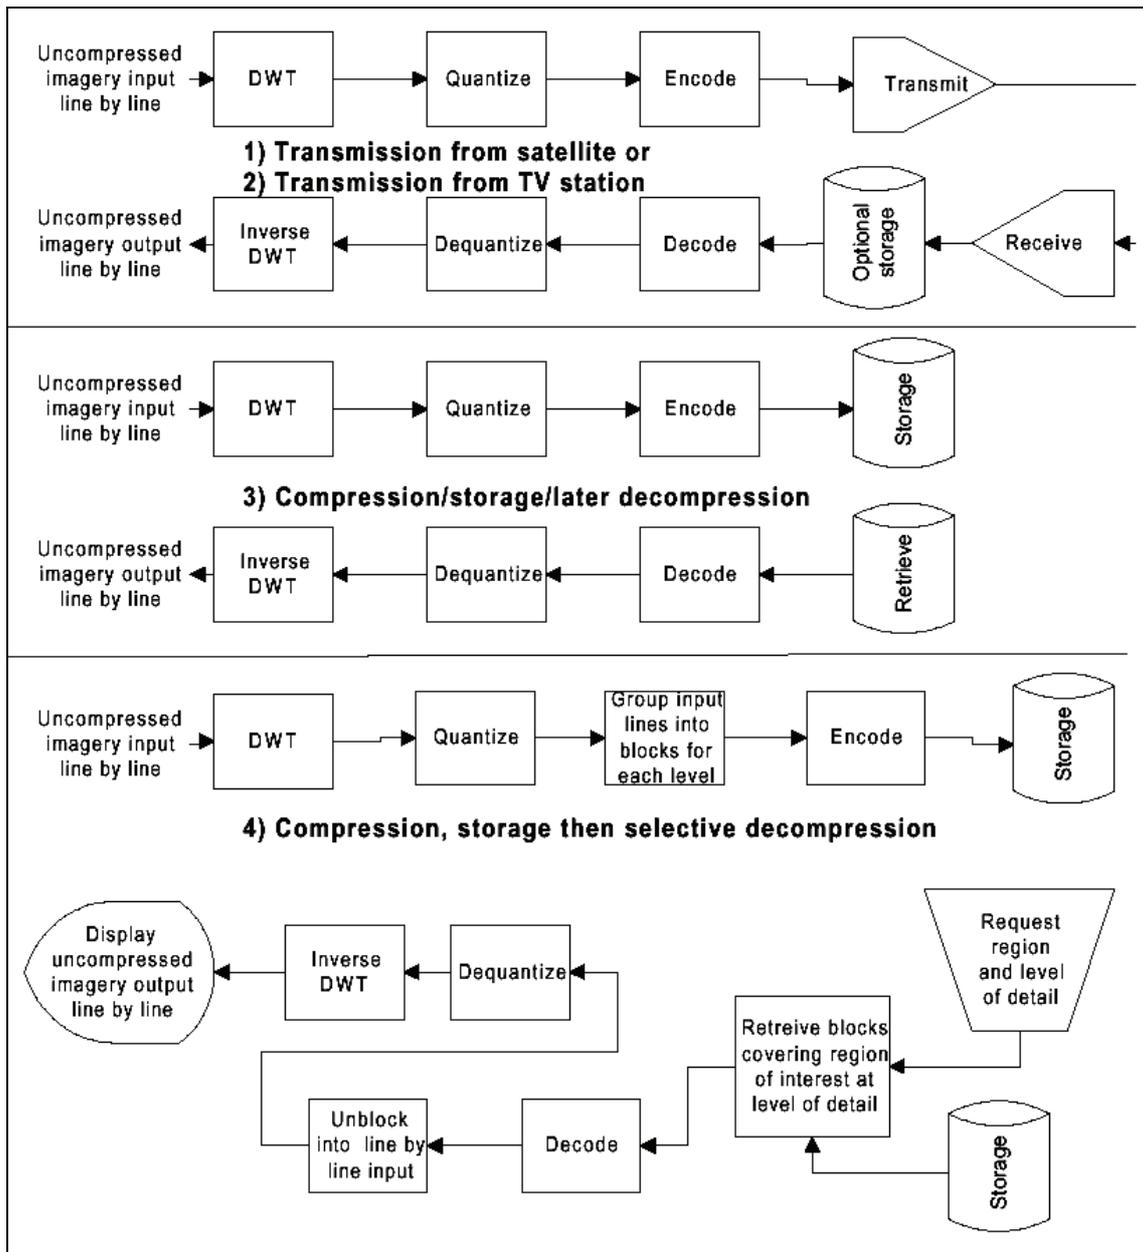


Figure 1: Different applications using the recursive DWT generation technology

2.3. ER Mapper Compressed Wavelet (ECW) versus other wavelet compression techniques

The ER Mapper ECW compression by Earth Resource Mapping uses advanced wavelet compression techniques.

Historically, the methods for doing Discrete Wavelet Transformations (DWT), which is an essential part of wavelet compression, have been memory based. This has limited the size of a file that can be compressed using wavelet compression.

Older existing compression techniques involve storing intermediate tile results on hard disk during compression and then later retrieving them to compute results for other tiles.

The ER Mapper ECW compression technique uses a patent pending recursive algorithm pipeline technique which does not require the use of disk storage while performing the DWT.

Most existing wavelet techniques enable compression of very large images - there is no inherent limitation to image size in the compression technology.

The primary advantages inherent in the ER Mapper ECW technique are that it is much faster, for several reasons:

- (a) The ECW technique does not require intermediate tiles to be stored to disk and then recalled during the DWT transformation and
- (b) The ECW technique takes advantage of CPU, L1 and L2 levels of cache to do its linear and unidirectional data flow through the DWT process.

Because the ER Mapper ECW method of DWT is much faster, this speed is used to provide more efficient compression in several ways:

- Multiple encoding techniques are used. Once an image has gone through DWT and quantization, it must be encoded. The ER Mapper ECW technique applies multiple, different, encoding techniques, and automatically chooses the best encoding method over each area of an image. Where multiple techniques are equally good, the method that is fastest to decode is chosen.
- Asymmetrical wavelet filters are used. Because of its speed, the ER Mapper ECW compression engine can use a larger and, therefore, slower DWT filter bank for DWT encoding. This enables the use of a smaller, faster inverse DWT filter for decoding. This means that decoding of imagery is much faster. A 15 tap floating point filter bank is used for DWT (compression), and a 3 tap integer based filter bank is used for the inverse DWT (decompression).
- Even with the additional processing carried out as detailed above, the ER Mapper ECW Compression Wizard is still 50% faster at compressing images than other compression techniques, when measured on the same file on the same computer.
- The decompression (which with most techniques can be for selective areas at selective levels of detail) is much faster for the ER Mapper ECW compressed image format than for other image formats, resulting in faster and more interactive use of compressed imagery.
- Because the ER Mapper ECW technique does not require intermediate DWT tiles to disk while performing the DWT step of image compression, it has other advantages that offer future benefits as well. These include:
 - Multiprocessor optimizations. The ER Mapper ECW DWT technique is inherently well suited for multi-CPU compression. Future versions of the ER Mapper ECW compression (and decompression) engine will take advantage of multi-CPU hardware for even faster compression/decompression.
 - Guaranteed latency. Because the ER Mapper ECW technique is a recursive algorithm pipeline technique, it has guaranteed latency (no disk IO is required) and guaranteed compression time with defined CPU performance. Although the ER Mapper ECW Compression Wizard reads uncompressed imagery from disk and writes compressed imagery to disk, this is simply an implementation and not an architecture requirement. The ECW technique could just as easily take a line by line stream of input uncompressed imagery, compress it, and emit a compressed stream of imagery. This means that the ER Mapper ECW compression/decompression technique can be applied to a range of applications which may have no local disk storage. These applications include transmission of HDTV signals, real time compression of imagery on satellites for reduced down-link data rate requirements, and compression of imagery on digital cameras.

Another significant difference between the ER Mapper ECW and other compression techniques is that the ER Mapper Compression Wizard is not a standalone compression product but is tightly integrated with other ER Mapper 6.0 functionality. It can, therefore, compress input from smart data algorithms as well as directly from uncompressed imagery.

This means that the other ER Mapper tools, such as the Ortho Wizard, the Mosaic Wizard, and the Balance Wizard, can be used to prepare seamless mosaics which can then be compressed using the Compression Wizard.

ER Mapper 6.0 includes a full suite of free plugins for GIS, CAD and Office applications, that enable those products to directly access compressed ECW image files. These free plugins may be freely distributed with your compressed data. There are no royalty fees for the free plugins, and there are no data royalty fees when compressing with the ER Mapper ECW format.

Software developers can license the ER Mapper ECW Decompressor free of charge, and use it within any application. The code is C based, and has a simple interface enabling it to be easily integrated into a range of products. Software developers can also license the ER Mapper Compression Engine. License fees depend on final usage.

3. COMPRESSION WITH THE ECW – COMPRESSOR

When you compress an image, you will be asked for the “Target Compression Ratio”. This is the desired compression ratio that you would like to achieve. For example, you might specify a ratio of 20:1 for an input file of 1,000MB to achieve a desired a 20MB compressed image (so the output image is 5% of the size of the input image).

After compressing the image, you may note that the actual compression rate achieved was in fact 40:1, resulting in an output file size of only 10MB in size.

These are the Target compression ratio (what you would like) and the Actual compression ratio (what was achieved). Except when compressing very small files (less than 2MB in size), the Actual compression ratio will generally be equal to or greater than the Target compression; sometimes significantly greater.

The reason for this is as follows:

When you specify a Target Compression ratio, the Compression Wizard uses this as a measure of how much information content to preserve in the image. If, however, your image has areas that are well suited to compression, a greater rate of compression may be achieved while still achieving the desired information content quality. The Compression Wizard uses multiple wavelet encoding techniques at the same time, and adapts and chooses the best technique depending on the area currently being compressed.

One example of this is an image that has large areas of water or desert. These can often be compressed with greater efficiency. Another example is a compressed image that consists of high-resolution airphotos, over lower resolution satellite imagery where there are no airphotos. Because the satellite images are lower spatial resolution, greater compression can be achieved in these areas of the image, while still preserving high quality detail in the airphoto area.

3.1. Typical compression ratios for different types of imagery and applications

The following table (Table 1) shows some typical target compression ratios for different applications. The suggested compression ratios reduce the amount of data to be stored in an effective way and provide the needed quality for the intended application.

Imagery	Application	Target Compression Ratio
Colour airphoto mosaic	High quality printed maps	25:1
Colour a20irphoto mosaic	Internet or email distribution	40:1
Greyscale airphoto mosaic	High quality printed maps	10:1 to 15:1
Greyscale airphoto mosaic	Internet or email distribution	15:1 to 30:1

Table 1: Suggested compression ratio for different applications

Some examples for the amount of data reduction are given in Table 2. As result of the ECW image compression you can store more than 80 colour airphotos on a single CD-ROM instead of three images of uncompressed size. So you can reduce costs effectively for production and delivery of your data. Also the access to the data speeds up very much not only because of the smaller amount of data but also because of the inherent access only of the level of detail which is needed for a request.

Type of imagery	Dimensions	Uncompressed size	Compression ratio	Compressed size
Greyscale airphoto	1 x 8,000 x 8,000	64MB	15:1	4MB
Colour airphoto	3 x 8,000 x 8,000	192MB	25:1	8MB
Scanned Topo map	3 x 40,000 x 60,000	7,200MB	50:1	144MB
Airphoto mosaic	3 x 50,000 x 100,000	15,000MB	25:1	600MB
State-wide mosaic	3 x 300,000 x 500,000	450GB	25:1	18GB

Table 2: Compressed sizes of different types of imagery

3.2. Results of different compression ratios

As an example from practice the following table (Table 1) shows some results for compressing the same image with different target compression ratios. The image is a colour airphoto with 10976 by 10914 pixel size, which was kindly given by Traveltainment AG, Aachen. The image was compressed with an AMD Athlon 750 PC with 256 MB RAM. The compression software was ER Mapper 6.1. The resulting compression ratios were smaller then the target compression ratios because of the kind of the environment covered by the airphoto. There are mostly shrubs with a big amount of textural information and no bigger areas with a similar surface. An overview of the airphoto is given in Figure 2.

Target compression ratio	Exact compression ratio	Compressed file size	Time needed for compression (m:s)	MB / second
10:1	9.1:1	37.3 MB	4:51	1.17
25:1	16.7:1	20.5 MB	4:02	1.41
50:1	30.1:1	11.4 MB	3:40	1.55
100:1	65.1:1	5.2 MB	3:31	1.62

Table 3: Results of compression of an colour airphoto with 10976 by 10914 pixel (original file size 342 MB) on an AMD Athlon 750 PC with 256 MB RAM, ER Mapper 6.1

Figure 3 to Figure 7 show a subset of the airphoto after the different compressions. Especially at the edges of roofs the loss of original information after decompression can be seen. Nevertheless also the compressed image with a target compression ratio of 50:1 gives good results for most applications.



Figure 2: Uncompressed aerial photograph (© Traveltainment)

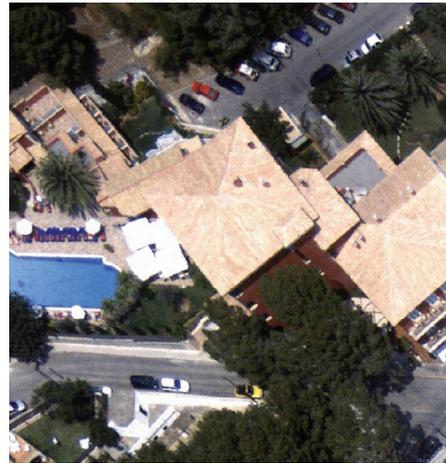


Figure 5: Target compression ratio 25:1, exact compression ratio 16.7:1 (© Traveltainment)

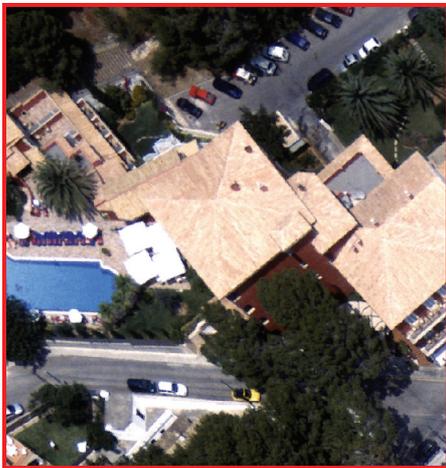


Figure 3: Detail of the uncompressed aerial photograph (© Traveltainment)

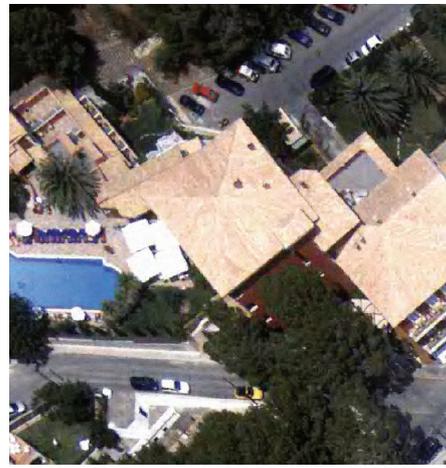


Figure 6: Target compression ratio 50:1, exact compression ratio 30.1:1 (© Traveltainment)

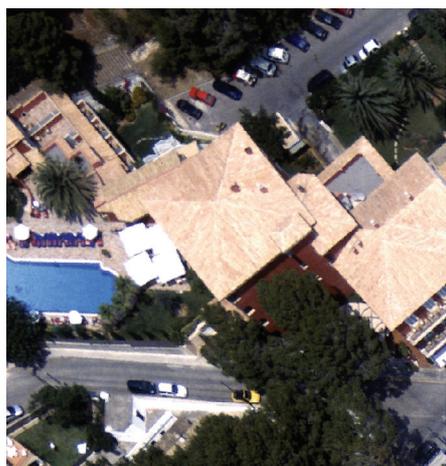


Figure 4: Target compression ratio 10:1, exact compression ratio 9.2:1 (© Traveltainment)



Figure 7: Target compression ratio 100:1, exact compression ratio 65.1:1 (© Traveltainment)

4. ACKNOWLEDGEMENTS

Many thanks to Traveltainment AG, Aachen, for the airphoto.

5. REFERENCES

Earth Resource Mapping (1999): Using and distributing wavelet compressed imagery, Internal technical white paper, unpublished