

Image compression

Digital Preservation guidance note 5

© Crown copyright 2008

You may re-use this information (excluding logos) free of charge in any format or medium, under the terms of the Open Government Licence. To view this licence, visit

nationalarchives.gov.uk/doc/open-government-licence or email psi@nationalarchives.gsi.gov.uk.

Where we have identified any third-party copyright information, you will need to obtain permission from the copyright holders concerned.

This publication is available for download at nationalarchives.gov.uk.

Contents

1	Introduction	3
2	Image compression considerations	3
2.1	Efficiency	3
2.2	Lossiness.....	4
3	Image compression algorithms.....	4
3.1	Run Length Encoding (RLE)	4
3.2	LZ Compressors	5
3.3	Huffman Encoding	5
3.4	Deflate	5
3.5	CCITT Group 3 and Group 4	5
3.6	Lempel-Ziff-Welch (LZW).....	6
3.7	JPEG	6
3.8	JPEG 2000	8
3.9	PNG compression.....	8
3.10	Fractal compression	8
3.11	HD Photo.....	9
4	Conclusions.....	10

1 Introduction

This document is one of a series of guidance notes produced by The National Archives, giving general advice on issues relating to the preservation and management of electronic records. It is intended for use by anyone involved in the creation of electronic records that may need to be preserved over the long term, as well as by those responsible for preservation.

This guidance provides advice on general issues, which should be considered by the creators and managers of electronic records when considering image compression. Please note that The National Archives does not specify or require the use of any particular image compression algorithm for records that are to be transferred. Choices should always be determined by the functional requirements of the record-creating process. Record creators should be aware however, of the implications of particular image compression techniques with regard to the long-term sustainability of digital images.

This guidance provides concise information about the most common image compression algorithms in use. It aims to help data creators and archivists to make informed choices about image compression issues. It should be read in conjunction with [Graphics file formats](#).

2 Image compression considerations

Computer graphics applications, particularly those generating digital photographs and other complex colour images, can generate very large file sizes. Issues of storage space, and the need to rapidly transmit image data across networks and over the Internet, has led to the development of a range of image compression techniques in order to reduce the physical size of files. Most compression techniques are independent of specific file formats - indeed, many formats support a number of different compression types. They are an essential part of digital image creation, use, and storage.

There are a number of factors to be considered when using compression algorithms:

2.1 Efficiency

Most algorithms are particularly suited to specific circumstances; these must be understood if they are to be used effectively. Some, for example, are more efficient at compressing monochrome images, whilst others yield better results with complex colour images.

2.2 Lossiness

Graphics compression algorithms fall into two categories:

-  Lossy compression achieves its effect at the cost of a loss in image quality, by removing some image information.
-  Lossless compression techniques reduce size whilst preserving all of the original image information, and therefore without degrading the quality of the image.

Lossy compression techniques should be treated with caution – if images are repeatedly migrated over time between different lossy formats, the image quality will be increasingly degraded at each stage. However, in some circumstances the use of lossy compression may be required, for example, to enable very large volumes of high-quality colour images to be managed economically. In such circumstances, visually-lossless compression should be used, which only removes image information which is invisible to the human eye at normal magnification.

Some compression algorithms are patented and may only be used under license. Others have been developed as open standards. This can be an important consideration in terms of both creation costs and long-term sustainability. The patenting of compression algorithms is a complex issue that is beyond the scope of this Guidance Note, though known issues are highlighted.

3 Image compression algorithms

This section describes the most commonly used compression algorithms for image data.

3.1 Run Length Encoding (RLE)

Run Length Encoding (RLE) is perhaps the simplest compression technique in common use. RLE algorithms are lossless. They work by searching for runs of bits, bytes, or pixels of the same value, and encoding the length and value of the run. As such, RLE achieves best results with images containing large areas of contiguous colour, and especially monochrome images. Complex colour images such as photographs, do not compress well and in some cases, RLE can actually increase the file size.

There are a number of RLE variants in common use; they are encountered in the TIFF, PCX and BMP graphics formats.

3.2 LZ Compressors

LZ compressors are a group of lossless compression schemes developed by Abraham Lempel and Jakob Ziv in 1977-1978.

- LZ77 compression underlies the Deflate algorithm (see 3.4), which is used in compressed archive formats such as PKZIP and the PNG compression algorithm (see 3.9)
- LZ78 compression is more commonly used for images, and forms the basis of the LZW algorithm (see 3.6)

3.3 Huffman Encoding

Developed by David Huffman in 1952, Huffman encoding is one of the oldest and most established compression algorithms. It is lossless and is used to provide a final compression stage in a number of more modern compression schemes, such as:

- JPEG (see 3.7) and Deflate (see 3.4)
- In modified form, it is also used in CCITT Group 3 compression (see 3.5)

3.4 Deflate

Deflate is a lossless algorithm based on LZ77 compression (see 3.2) and Huffman encoding (see 3.3). It was developed by Phil Katz in 1996 for use in the PKZIP compressed archive format, and forms the basis for PNG compression (see 3.9).

3.5 CCITT Group 3 and Group 4

Officially known as CCITT T.4, Group 3 is a compression algorithm developed by the International Telegraph and Telephone Consultative Committee in 1985 for encoding and compressing 1-bit (monochrome) image data. Its primary use has been in fax transmission, and it is optimised for scanned printed or handwritten documents.

Group 3 is a lossless algorithm, of which two forms exist: one-dimensional (which is a modified version of Huffman encoding) and two-dimensional, which offers superior compression rates. Due to its origin as a data transmission protocol, Group 3 encoding incorporates error detection codes.

Group 4 compression, officially CCITT T.6 is a development of the two-dimensional Group 3 standard, which is faster and offers compression rates which are typically double those of Group 3. Like Group 3, it is lossless and designed for 1-bit images.

However, being designed as a storage rather than transmission format, it does not incorporate the error detection and correction functions of Group 3 compression.

Group 3 and 4 compression are most commonly used in the TIFF file format.

 The full technical specifications for Group 3 and 4 compression are published in CCITT Blue Book, 1989, Volume VII, Fascicle VII.3: Terminal equipment and protocols for telematic services, recommendations T.0 – T.63.

3.6 Lempel-Ziff-Welch (LZW)

The Lempel-Ziff-Welch compression algorithm was developed by Terry Welch in 1984, as a modification of the LZ78 compressor (see 3.2). It is a lossless technique that can be applied to almost any type of data, but is most commonly used for image compression. LZW compression is effective on images with colour depths from 1-bit (monochrome) to 24-bit (True Colour).

The patent for the LZW algorithm is owned by Unisys Corporation, which has licensed its use in a variety of file formats, most notably CompuServe's GIF format. (See [Graphics file formats](#) for more information on GIF). The licensing applies to implementations of the LZW algorithm, and not to individual files which utilise it. The US patent expired in June 2003 and the UK patent expired in June 2004.

LZW compression is encountered in a range of common graphics file formats, including TIFF and GIF.

 The definitive source for information on LZW compression is Welch, T A, 1984, A technique for high performance data compression, IEEE Computer, 17: 6.

3.7 JPEG

The JPEG compression algorithm has its origins in moves to develop compression techniques for the transmission of colour and greyscale images. Developed in 1990 by the Joint Photographic Experts Group of the International Standards Organisation (ISO) and CCITT,

JPEG is a lossy technique that provides best compression rates with complex 24-bit (True Colour) images. It achieves its effect by discarding image data that is imperceptible to the human eye, using a technique called Discrete Cosine Transform (DCT). It then applies Huffman encoding to achieve further compression.

JPEG comprises a baseline specification, for which a number of optional extensions have been defined, including:

-  Progressive JPEG allows a JPEG decoder to build up and display an image progressively, rather than waiting for all the image data to be received, and can be useful for applications which need to stream image data.
-  Arithmetic encoding is an extension of baseline JPEG which offers higher compression rates, but is slower and subject to patent.
-  Lossless JPEG uses a different algorithm to baseline JPEG to provide lossless compression. Support for lossless JPEG is very limited.

The JPEG specification allows users to set the degree of compression, using an abstract Quality Setting. This provides a trade-off between compression rate and image quality: the higher the setting, the better the quality of the resultant image but at the cost of a larger file size. It is important to note that the Quality Setting is not an absolute value: different JPEG encoders use different scales, although 1 to 100 is typical. In addition, with some implementations, 100 would represent maximum compression rather than maximum image quality. It is also essential to be aware that even the maximum quality setting for baseline JPEG involves some degree of lossy compression. Repeated saving of an image will always lead to increasing degradation of the image quality.

The arithmetic encoding extension was subject to patent, but this expired in 2004.

JPEG compression is used in the JPEG File Interchange Format (JFIF), SPIFF and TIFF.

-  The full technical specifications for baseline JPEG compression have been published as an international standard (ISO/IEC 10918 Part 1). Extensions to baseline JPEG are described in ISO/IEC 10918 Part 3.

3.8 JPEG 2000

JPEG 2000 is a replacement for the JPEG algorithm, developed by the ISO JPEG group in 2000. It provides for lossy and lossless compression, and uses wavelet compression to achieve higher compression rates with a lower corresponding reduction in image quality.

JPEG 2000 may utilise some patented technologies, but is intended to be made available on a license- and royalty-free basis.

The JPEG 2000 standard defines a minimum file interchange format (JP2), in a similar manner to JFIF and SPIFF. Support for JPEG 2000 is now available in many commercial software packages.

 The full technical specifications for JPEG 2000 compression have been published as an international standard (ISO/IEC 15444 Part 1).

3.9 PNG compression

PNG compression was developed in 1996 as part of the PNG file format (see [Graphics file formats](#)), to provide a non-proprietary alternative to the LZW compression employed by GIF and other file formats (see 3.6).

PNG compression uses the Deflate compression method (see 3.4). It is a lossless algorithm and is effective with colour depths from 1-bit (monochrome) to 48-bit (True Colour).

PNG compression is unencumbered by patent and free to use. It is implemented only in the PNG file format.

 The full technical specifications for PNG compression have been published as part of the full format specification, in RFC-2083, and as a W3C recommendation. Version 1.2 was released as an ISO standard (ISO/IEC International Standard 15948) in 2004.

3.10 Fractal compression

Fractal compression uses the mathematical principles of fractal geometry to identify redundant repeating patterns within images. These matching patterns may be identified through performing geometrical transformations, such as scaling and rotating, on elements of the image. Once

identified, a repeating pattern need only be stored once, together with the information on its locations within the image and the required transformations in each case. Fractal compression is extremely computationally intensive, although decompression is much faster. It is a lossy technique, which can achieve large compression rates. Unlike other lossy methods, higher compression does not result in pixelation of the image and, although information is still lost, this tends to be less noticeable. Fractal compression works best with complex images and high colour depths.

The original fractal compression algorithm was developed by Michael Barnsley in 1991. However, the algorithm is patented and supported by few commercial products. It is not implemented in any common graphics file formats.

3.11 HD Photo

HD Photo compression is a compression algorithm introduced by Microsoft in its HD Photo format (previously Windows Media Photo). This supports lossy and lossless compression with claimed performance improvements over JPEG 2000. The algorithm is patented, but Microsoft has announced that the specification will be released under their Open Specification Promise. Support for HD Photo is beginning to grow amongst third-party software products.

HD Photo is under consideration as an ISO standard by the Joint Photographic Experts Group, tentatively known as JPEG XR.



The technical specifications for HD Photo are available from the [Microsoft website](#).

4 Conclusions

The table below summarises the lossiness of the algorithms described, and the circumstances in which they are most efficient:

Algorithm	Lossiness	Efficient with
RLE	Lossless	Monochrome or images with large blocks of colour
LZ Compressors	Lossless	All images
Huffman Encoding	Lossless	All images
Deflate	Lossless	All images
CCITT Group 3 & 4	Lossless	Monochrome images
LZW	Lossless	All images
JPEG	Lossy (lossless extension available)	Complex, True Colour images
JPEG 2000	Lossy, lossless supported	Complex, True Colour images
PNG	Lossless	All images
Fractal	Lossy	Complex, True Colour images
HD Photo	Lossy, lossless supported	Complex, True Colour images

It is recommended that algorithms should only be used in the circumstances for which they are most efficient. It is also recommended that archival master versions of images should only be created and stored using **lossless** algorithms, or **visually lossless** levels of lossy compression.

The Intellectual Property Rights status of a compression algorithm is primarily an issue for developers of format specifications, and software encoders/decoders. However, the use of open, non-proprietary compression techniques is recommended for the purposes of sustainability.