

Fpack and Funpack Utilities for FITS Image Compression and Uncompression

William Pence, NASA Goddard Space Flight Center, Greenbelt, MD 20771

Rob Seaman, National Optical Astronomy Observatory, Tucson, AZ 85719

Rick White, Space Telescope Science Institute, Baltimore, MD 21218

[DRAFT] 13 June 2008

1. Introduction

Fpack is a utility program for optimally compressing images in the FITS (Flexible Image Transport System) data format (see <http://fits.gsfc.nasa.gov>). The associated funpack program restores the compressed image file back to its original state (as long as a lossless compression algorithm is used). These programs may be run from the host operating system command line and are analogous to the gzip and gunzip utility programs except that they are optimized for FITS format images and offer a wider choice of compression algorithms.

Fpack stores the compressed image using the FITS tiled image compression convention (see http://fits.gsfc.nasa.gov/fits_registry.html). Under this convention, the image is first divided into a user-configurable grid of rectangular tiles, and then each tile is individually compressed and stored in a variable-length array column in a FITS binary table. By default, fpack usually adopts a row-by-row tiling pattern. The FITS image header keywords remain uncompressed for fast access by FITS reading and writing software.

The tiled image compression convention can in principle support any number of different compression algorithms. The fpack and funpack utilities call on routines in the CFITSIO library (<http://heasarc.gsfc.nasa.gov/fitsio>) to perform the actual compression and uncompression of the FITS images, which currently supports the GZIP, Rice, H-compress, and the PLIO IRAF pixel list compression algorithms.

The fpack and funpack utilities were originally designed and written by Rob Seaman (NOAO). William Pence (NASA) added further enhancements to the utilities and to the image compression algorithms in the underlying CFITSIO library. Rick White (STScI) wrote the code for the Rice and Hcompress algorithms, and provided general advice on other image compression issues.

2. Benefits of fpack

Using fpack to compress FITS images offers a number of advantages over the other commonly used technique of externally compressing the whole FITS file with gzip:

1. fpack generally offers higher compression ratios and faster compression and uncompression

- speed than gzip.
2. The FITS image header keywords remain uncompressed and thus can be read or written without any additional overhead.
 3. Each HDU of a multi-extension FITS file is compressed separately, thus it is not necessary to uncompress the entire file to read a single image in a multi-extension file.
 4. The capability of dividing the image up into tiles before compression enables faster access to small subsections of the image because only those tiles contained in the subsection need be uncompressed.
 5. The compressed image is itself a valid FITS file and thus can be manipulated by other general FITS utility software.
 6. Fpack supports lossy compression techniques that can achieve significantly higher compression factors than the lossless compression algorithms in situations where it is not necessary to exactly preserve every bit of the original image pixel values. This is especially relevant when compressing 32-bit floating point FITS images for which there is often little justification for preserving the full numerical precision (6 - 7 decimal places) of each pixel value.
 7. Fpack and Funpack automatically update the CHECKSUM keywords in the compressed and uncompressed files to help verify the integrity of the FITS files.
 8. Software applications that are built on top of FITS access libraries such as CFITSIO, that internally support the tiled image compression technique, are able to directly read and write the FITS images in their compressed form, thus reducing the amount of disk storage space needed by users.

3. General fpack usage guidelines

In its simplest application, fpack can be used to minimize the data storage and network bandwidth resources needed to archive and distribute FITS images. In this scenario, the data producer compresses the FITS image files with fpack before placing them in the data archive; after downloading the compressed files, each end user of the images then runs funpack to convert them back into the standard FITS image format before doing any further analysis.

The benefit of using fpack to compress FITS images is much greater, however, when the user's analysis software is capable of directly reading and writing the files in the compressed form. This reduces both the required amount of local disk space and the time needed to copy the files from one location to another. Directly reading or writing FITS images in the compressed format may require more CPU resources, but this is offset by a gain in I/O performance because fewer bytes of data need to read or written.

Any software application that uses the CFITSIO library (<http://heasarc.gsfc.nasa.gov/fitsio>) to read and write FITS images will transparently inherit the ability to read or write tile-compressed images. The image compression or uncompression is performed by the CFITSIO library routines, so in general, the application program itself does not need to know anything about the tiled compressed image format. The main exception to this is that when writing compressed images, the application program may need

to call an additional CFITSIO routine to define which compression algorithm to use, along with the values of other optional compression parameters. The fpack and funpack utilities are themselves examples of applications that use CFITSIO to perform the compression and uncompression operations on the images.

Besides CFITSIO, the IRAF data analysis system also provides some support for the tile-compressed image format. As this format become more widely used in the future, it is anticipated that other analysis systems will also add support for this tiled image compression format. In the meantime, however, users have the option of using funpack to uncompress the images back into standard FITS images that should be compatible with other astronomical image analysis software that does not directly support the compressed format.

4. Compression versus noise

When dealing with lossless compression of images with integer-valued pixels, the amount of compression that will be achieved depends almost completely on one simple factor: the amount of the random noise that is present in the pixel values in the image. As the amount of noise is reduced, the compression ratio increases exponentially. By comparison, other image attributes, such as the mean pixel value, or the presence of image structures such as stars, spectral lines, or large scale gradients across the image, have little or no effect on the compression factor.

In order to quantify this effect, fpack uses a robust algorithm for measuring the amount of noise in "background" areas of the image. This algorithm was originally developed to measure the signal-to-noise in spectroscopic data (see the June 2007 ST-ECF Newsletter). In particular, we use the 3rd order "median absolute difference" formula to compute the standard deviation of the (assumed) Gaussian distributed noise for the pixels in each row of the image:

$$\text{standard deviation} = 0.6052 * \text{median}(\text{abs}(-x(i-2) + 2x(i) - x(i+2))))$$

where $x(i-2)$ is the value of the pixel 2 spaces to the left of pixel i , and $x(i+2)$ is the value of the pixel 2 spaces to the right, and the median value is computed over all the pixels in each row of the image. This formula is much less sensitive to outlying pixel values than the usual formula for computing the standard deviation of a sample of values. The overall noise value for the image as a whole is then computed from the mean of the median values for each row of the image.

Extensive testing on synthetic images with known amounts of Gaussian noise, as well as on large sets of real astronomical images, has shown that the amount of compression that is achieved can be accurately predicted from the noise value that is computed with the above formula. Table 1 shows the typical compression factor that is achieved, when using the default Rice compression algorithm, with images containing varying amounts of noise:

Table 1.

noise	Compression Ratio
3	4.0
5	3.0
10	2.5
30	2.0
100	1.7

Since the Poisson noise in an ideal photon-counting detector is equal to the square-root of the number of photons, the measured noise in a image typically increases as the mean pixel value increases. This also means that the noise level also increases with the exposure time, all other factors remaining constant. To take the NOAO Mosaic CCD camera as an example, "bias" frame images that have 0 exposure time have a very low mean pixel level and consequently have low noise levels in the range 2 - 5. These bias frame images compress very well. Typical images of the night sky, on the other hand, have significantly higher mean count levels, and have typical noise levels in the range 20 - 50. The "flat field" images intentionally have very high mean pixel values and consequently have the highest noise values, typically in the range 80 - 200. Thus there is a natural progression in that the shortest exposure images with low count levels compress the most, and the longest exposure or highest count level images compress the least.

These same rules apply when compressing 32-bit integer images, only the compression factor that is achieved is twice that of a 16-bit image, given the same noise level. If a 16-bit image and a 32-bit image have the same noise level, then the compressed images will have almost exactly the same size. Since the 32-bit image was originally twice as big as the 16-bit image (assuming the same image dimensions), then the compression ratio for the 32-bit image will be twice that of the 16-bit image..

5. Specific fpack compression guidelines

The optimal compression technique to use with fpack depends on the data type of the FITS image (either integer or floating point) and on whether lossless or lossy compression is to be performed. The following subsections provide guidelines on using fpack in each of these cases.

5.1 Lossless compression of integer FITS images

When compressing integer images which have BITPIX = 8, 16, or 32, the default Rice compression algorithm and row-by-row tiling pattern generally provides the optimum compression. One can experiment using different compression algorithms (-g for GZIP, -h for Hcompress, and -p for the IRAF PLIO algorithm) or different tiling patterns (with the -t option) to test how these affect the compression

ratio and speed. The `fpack` program provides the `-T` test option that is convenient for comparing the compression ratio and the amount of time needed to compress and uncompress the image when using each of the 3 main compression algorithms (Rice, Hcompress, and GZIP).

5.2 Lossy compression of integer images

The H-compress algorithm may be used in a lossy mode to achieve higher compression of integer images. This mode will obviously not exactly preserve the pixel values in the original image, but this may be appropriate if the image is noisy, or if the image is intended mainly for qualitative purposes. To use this option in `fpack`, choose the H-compress algorithm, and set the `-s` option to a positive value, typically somewhere in the range 1.0 to 100. Larger values will produce greater amounts of compression. Note that it may be necessary to increase the default tile size when doing lossy compression, otherwise the boundaries of the tiles may become visible in the compressed image.

5.3 Lossy compression of floating point FITS images

When compressing floating point FITS images which have `BITPIX = -32` or `-64`, `fpack` always converts the pixel values into 32-bit integers using a linear scaling function:

$$\text{integer_value} = (\text{floating_point_value} - \text{ZERO_POINT}) / \text{SCALE_FACTOR}$$

This array of scaled integers is then compressed using the specified compression algorithm. When the image is subsequently uncompressed, the integer values are inverse scaled to closely, but not exactly, reproduce the original floating point pixel values. Separate scale and zero point values are computed separately for each tile of the image. This scaling of floating point images is performed because a) many compression algorithms can only operate on integer arrays, and b) even if the algorithm can compress floating point arrays (e. g., GZIP), the amount of compression that is achieved is generally poor as compared to compressing the equivalent array of scaled integers.

The value of `SCALE_FACTOR` in the scaling function controls how closely the inverse scaled values approximate the original floating point values: decreasing `SCALE_FACTOR` reduces the spacing between the quantized levels in the inverse-scaled values and thus more closely reproduces the original pixel values. However, this also magnifies the dynamic range and the noise level in the integer array that is to be compressed which, as discussed in section 4, adversely affects the amount of compression that is achieved. Thus, there is a direct trade-off between providing more precision (by decreasing `SCALE_FACTOR`) or achieving a greater amount of compression (by increasing `SCALE_FACTOR`).

It is not easy to directly determine an appropriate `SCALE_FACTOR` value to use with a given image, therefore `fpack` provides instead a quantization parameter called "q" for specifying how closely the inverse-scaled integer pixel values must approximate the original floating point pixel values, relative to the measured noise in background areas in the image. With the default value of `q = 16`, the image pixel values will be quantized so that the spacing between the adjacent discrete levels is equal to

the measured R.M.S. noise in the background regions of the image divided by 16. In other words, the pixel values are recorded with a factor of 16 times more accuracy than the measured background noise level. For example, if the RMS noise in a tile of floating point image has a value of 25.0, then the pixel values in the compressed image will be quantized into levels that are separated by $25 / 16 = 1.56$. The maximum difference between the pixel values in the compressed image and the original image will be half this value. Increasing the value of q will produce compressed images that more closely approximate the pixel values in the original floating point image, but will also increase the size of the compressed image file. Further details of this floating point compression scheme are given in the ADASS paper by White and Greenfield, 1999.

Once the floating-point pixel values have been converted to scaled integers, they may be compressed using either the Rice, GZIP, or H-compress algorithms. In most cases the default Rice algorithm provides the best compromise between speed and compression factor. Rice is generally faster and provides better compression than GZIP. While H-compress can often produce slightly better compression ratios than the Rice algorithm, this small benefit is usually not worth the much larger CPU times that are required to compress or uncompress the image using H-compress.

The q parameter directly determines the amount of compression in floating point images because the noise in the scaled integer array is numerically equal to q (so for example, if $q = 16$ then the noise in the scaled integer image will also equal 16). Table 2 shows the approximate compression ratio that will be achieved for floating point images as a function of q when using the default Rice compression algorithm. As was noted in section 4. these compression ratios are approximately twice those shown in Table 1 for 16-bit integer images with the equivalent noise level.

Table 2.

q	Compression Ratio
4	6.0
8	5.1
16	4.4
32	3.9
64	3.5

It must be emphasized that each fpack user is ultimately responsible for determining the appropriate q value to use when compressing floating point images; using too small a value of q could result in unrecoverable loss of some of the information that was present in the original floating point image. This loss may be acceptable in cases where the image is only used for qualitative purposes. Using too large a value of q , on the other hand, will simply preserve more of the random noise in the image and reduce the amount of compression that is achieved. Anecdotally, tests performed at the Space Telescope Institute and elsewhere using $q = 16$ did not detect any significant difference between various

photometric or astrometric quantities in a sample of compressed astronomical images, as compared with the same quantities derived directly from the original uncompressed image. This is no guarantee, however, that the default value of $q = 16$ is suitable for all applications, so users are strongly urged to perform quantitative tests using different values of q to determine the appropriate level to use for their particular application.

6. fpack command-line parameters

The fpack program is invoked on the command line like other host-level utility programs:

fpack [OPTION]... [FILE]...

The specified options must appear before the list of files to be compressed. The file names may contain the usual wildcard characters that will be expanded by the Unix shell.

Compression algorithm: select 1 of the following options:

- r** Rice [default], or
- h** Hcompress, or
- g** GZIP (per-tile), or
- p** IRAF pixel list compression algorithm. This can only be applied to images whose pixel values all lie in the range 0 to 2^{24} (16777216). Note that the PLIO algorithm cannot be used with FITS images that have unsigned integer values because the pixel values in these images are internally shifted into the range of negative signed integers.
- d** no compression (debugging mode)

Tiling pattern specification:

When using the Rice, GZIP, or PLIO compression algorithms, the default tiles each contains 1 row of the image (i.e., the tiles are one dimension and contain NAXIS1 pixels). The Hcompress algorithm requires that the tiles be 2-dimensional, therefore the default is to use 16 rows of the image per tile. If this would cause the last partially full tile of the image to only contain a small number of rows, then a slightly different tile size is chosen so that the last tile is more equal in size.

The default tile sizes can be overridden with one of the following fpack options:

- w** compress the whole image as a single large tile
- t <axes>** comma separated list of tile dimensions (e. g., -t 200,200 will produce tiles that are 200 x 200 pixels in size)

Compression parameters for floating point images:

-q <level> Quantization level when compressing floating point images. See the section 5.3 on compressing floating point images for a description of this parameter. It is important to realize that fpack does not exactly preserving the original pixel values when compressing floating point images. Users should carefully evaluate the compressed images (e.g., by uncompressing them with funpack) to make sure that any essential information in the original image has not been lost.

Positive q values are interpreted as relative to the R.M.S. noise in the image derived using the formula described in section 4. In some instances it may be desirable to specify the exact q value (not relative to the measured noise), so that all the tiles in the image, and all the images in a dataset, are compressed using the identical value, regardless of slight variations in the measured noise level. This can be done by specifying the negative of the desired value. The -T option can be used to calculate the noise level in the image.

-n <noise> Rescale the pixel values in a previously scaled image to improve the compression ratio by reducing the R.M.S. noise in the image. This option is intended for use with images that use scaled integers to represent floating point pixel values, and in which the scaling was chosen so that the range of the scaled integer values covers the entire allowed range for that integer data type (e.g., -32768 to +32767 for 16-bit integers and -2147483648 to +2147483647 for 32-bit integers). The measured R.M.S. noise in these integer images is typically so huge that they cannot be effectively compressed (because the compression ratio depends directly on the noise). This -n option rescales the pixel values so that the R.M.S. noise will be equal to the specified value. Appropriate values of n will likely be in the range from 8 (for low precision and the high compression) to 64 (for the high precision and lower compression). Users should read the section on compressing floating point images, above, for guidelines on choosing an appropriate value for n that does not lose significant information in the image.

Parameters for lossy compression of integer images:

-s <scale> Scale factor for lossy compression when using Hcompress. The default value is 0 which implies lossless compression. Positive scale values are interpreted as relative to the R.M.S. noise in the image. Scale values of 1.0, 4.0, and 10.0 will typically produce compression factors of about 4, 10, and 25, respectively, when applied to 16-bit integer images. In some instances it may be desirable to specify the exact scale value (not relative to the measured noise), so that all the tiles in the image, and all the images in a dataset, are compressed with the identical scale value, regardless of slight variations in the measured noise level. This is done by specifying the negative of the desired value (e.g. -30., which would be equivalent to specifying a scale value of 2.0 in an image that has RMS noise = 15.).

It is important to realize that this option achieves the high compression ratios at the expense of not exactly preserving the original pixel values in the image. Users should carefully evaluate the compressed images (e.g., by uncompressing them with funpack) to make sure that any essential information in the image has not been lost.

Output file naming parameters:

- F** force the input file to be overwritten by the compressed file with the same name. This is only allowed when a lossless compression algorithm is used.
- P <pre>** prepend the <pre> string to the input file name to generate the name of the compressed output file.
- A <suffix>** append <suffix> to the root of the input file name to generate the name of the compressed output file.
- S** output the compressed FITS file to the STDOUT stream (to be piped to another task)

Other miscellaneous parameters:

- v** verbose mode; list each file as it is processed
- T** produce a report that compares the compression ratio and the compression and uncompression times for each of the main compression algorithms. The input file remains unchanged and is not compressed
- O <filename>** write the comparison test report (produced by the -T option) to a file in a format that is suitable for ingest into a data base for further analysis.
- L** list all the extensions in all the input, files. No compression is performed.
- C** don't update FITS checksum keywords
- H** display a summary help file that describes the available fpack options
- V** display the program version number

7. funpack command-line parameters

funpack shares many of the same parameters as fpack as shown below:

Output file naming parameters:

- F** force the input file to be overwritten by the uncompressed file with the same name. This is only allowed when a lossless compression algorithm is used.
- P <pre>** prepend the <pre> string to the input file name to generate the name of the uncompressed output file.
- A <suffix>** append <suffix> to the root of the input file name to generate the name of the uncompressed output file.

-S output the uncompressed FITS file to the STDOUT stream (to be piped to another task)

Other miscellaneous parameters:

- v** verbose mode; list each file as it is processed
- L** list all the extensions in all the input, files. No uncompression is performed.
- C** don't update FITS checksum keywords
- H** display a summary help file that describes the available funpack options
- V** display the program version number

8. Building fpack and funpack

The latest versions of the fpack and funpack C source code are available from <http://heasarc.gsfc.nasa.gov/fitsio/fpack>. These programs are also included in the CFITSIO source file distributions (but not necessarily the latest version) available at <http://heasarc.gsfc.nasa.gov/fitsio>.

To build the software on unix systems, first download and build the CFITSIO library. If necessary, untar the latest version of the fpack and funpack source code into the CFITSIO directory, overwriting the older version. Then enter the commands

```
make fpack
```

```
make funpack
```

in that directory. This will create the fpack and funpack executable files which may be copied to any other suitable directory (e.g. the local /bin directory).