5.5.1 LOSSY COMPRESSION

As reported in Reference 5.5.1-1 *(excluding “Data and Software” section)*

Reference 5.5.1-1 - Web Page Report, http://astrosun.tn.cornell.edu/faculty/squyres/compression/compression.html, David O’brien (Cornell University)

5.5.1.1 DESCRIPTION OF SOFTWARE AND COMPRESSION PARAMETERS

5.5.1.1.1 Compression Software

The compression/decompression software used was provided by the Matra Marconi company in France, and has the same algorithm and features as the hardware compressor which they designed for the Cassini mission. It is based on the JPEG algorithm (see below), and has numerous variable parameters which affect the compression level and compression efficiency. In the actual hardware compressor, there are 4 selectable `pages' in memory, called PMEM pages (Parameter MEMory), which contain compression parameters optimized for different image types. In the software compressor, these `pages' are stored in a data file. There is also a parameter called the `B value', which controls the degree to which an image is compressed, and the `GOB Length' which controls the amount of error correction coding implanted in the compressed data stream. Other parameters which may be set are the NLS bit, which controls the order of certain steps in the compression process, and the optimized decompression option in the decompression program, which provides better reconstruction of the compressed image than the standard decompression scheme. These parameters are described in fuller detail in the section on compression parameters.

5.5.1.1.2 The JPEG Algorithm

The JPEG algorithm is based on the fact that the human eye is more sensitive to slow variations in intensity than rapidly varying ones, i.e. the eye is more sensitive to low spatial frequencies than high spatial frequencies. Furthermore, in most natural scenes (including planets and moons), the majority of visual information is in the form of low spatial frequency components. JPEG operates by throwing out some of the high spatial frequency components in an image, thus decreasing the amount of information that must be coded for while keeping most of the visually important information intact.

The JPEG algorithm can be outlined as follows:

1) The 8 bit image to be compressed is broken into 8 by 8 pixel blocks. Each pixel is `level shifted' by subtracting 128 from it.

2) The blocks are transformed by the Discrete Cosine Transform (DCT). This is an analogue of the Fourier transform, and essentially breaks each block into its spatial frequency components.

3) The high spatial frequency components are reduced by dividing each transformed 8 by 8 block by a `quantization matrix'. A `scale factor' is used to determine the degree to which high spatial frequency components are attenuated.

4) The information remaining in the blocks is then Huffman coded, formatted, and transmitted or written to a file. Decompression is performed by reversing the compression process.

5.5.1.1.3 Compression Parameters

5.5.1.1.3.1 PMEM Page

Each PMEM page contains a quantization matrix and a table of scale factors, as well as a Huffman coding table optimized for a certain type of image (busy, sky, or atmospheric images). Additionally, the first PMEM page (page 0) contains information which controls the format of the picture header and the error control headers. In pages 0, 2, and 3 the quantization matrix is flat (each element = 255), while in page 1 it is scaled to discard more high frequency information. Furthermore, the table of scale factors is the same in pages 0, 1, and 3 but is different for page 2. PMEM pages are selected by setting the values of two bits--Malgo and TB. The pages are summarized below:

Malgo	**TB**	**Page**	**Quantization Matrix**	**Huffman Table**
0	0	0	Flat (255)	Busy
1	0	1	Scaled	Busy
0	1	2	Flat (255)	Sky
1	1	3	Flat (255)	Atmospheric

5.5.1.1.3.2 B Value

The B value ranges from 0 to 15, and selects the value from the scale factor table to be used. High B values give higher compression. It should be noted that since PMEM page 2 has a different scale factor table than the other pages, a given B value generally selects a different scale factor when used with page 2 than it does with the other pages.

5.5.1.1.3.3 GOB Length

GOB length is a parameter ranging from 0 to 255, and determines how frequently error correction headers will be placed in the data stream. These headers allow an image to be mostly recovered if there is an error in transmission. If GOB length is 0, there will be no headers. For other values, a GOB length L will cause error correction headers to be placed after each group of L image blocks. If a GOB length of 0 is used, no protection is given against transmission errors. For a GOB length L, at most L image blocks will be lost from a single error. However, for nonzero L values, the smaller L is, the more additional information is included in the data stream--better error protection is balanced by a decrease in compression ratio.

5.5.1.1.3.4 NLS Bit

The JPEG algorithm outlined above specifies a level shift before the DCT. The Matra compressor allows this level shift to occur after the DCT in the actual compressor/formatter (COFO) section of the hardware or software, if desired. A NLS (No Level Shift) value of 1 causes the level shift to occur before DCT, as in the standard JPEG algorithm, while a NLS value of 0 causes the level shift to occur in the COFO. Neither option changes the performance of the compressor, but having the level shift occur in the COFO (NLS=0) is slightly easier for the hardware implementation of the compressor.

5.5.1.1.3.5 Optimized Decompression

This option in the decompression software constructs a histogram of the coded values and uses an interpolation method on it to provide better reconstruction of the coded data.

5.5.1.2 IMAGE ENTROPY

5.5.1.2.1 Description

Image entropy is a quantity which is used to describe the `business' of an image, i.e. the amount of information which must be coded for by a compression algorithm. Low entropy images, such as those containing a lot of black sky, have very little contrast and large runs of pixels with the same or similar DN values. An image that is perfectly flat will have an entropy of zero. Consequently, they can be compressed to a relatively small size. On the other hand, high entropy images such as an image of heavily cratered areas on the moon have a great deal of contrast from one pixel to the next and consequently cannot be compressed as much as low entropy images.

5.5.1.2.2 Calculating Image Entropy

Image entropy as used in my compression tests is calculated with the same formula used by the Galileo Imaging Team:

$$Entropy = -\sum_{i} P_i Log_2 P_i$$

In the above expression, P_i is the probability that the difference between 2 adjacent pixels is equal to i, and Log_2 is the base 2 logarithm.

Below is an IDL program I wrote to calculate the entropy of an image using this expression.

```
function entropy_func, arrayin

;-------------------------------------------------------------------

; Modified from the procedure to return the entropy value

; Calculates the line-by-line image entropy of the image stored in
; 'arrayin'. Prints the entropy value to screen.  If the array is of
; integer type, a 12 bit image is assumed.  If it is of byte type,
; an 8 bit image is assumed.

; Usage: entropy_value = entropy(image_array)

; David O'Brien, 5/30/97 -- problem fixed 8/1/97

;-------------------------------------------------------------------

; Get array dimensions and type

dimensions = size(arrayin)
xsize = dimensions(1)
ysize = dimensions(2)
type = dimensions(3)

; Fix array to integer type

array = fix(arrayin)
```

```
; Find correct size for probability array and create it

if type eq 1 then array_size = 256*2-1
if type eq 2 then array_size = 4096*2-1
if (type ne 1) and (type ne 2) then print, 'Not a byte or integer array'

prob_array = fltarr(array_size)

; Loop over image array elements and count occurences of each possible
; pixel to pixel difference value.  Store these values in prob_array

for j = 0, ysize-1 do $

        for i = 0, xsize-2 do begin

                diff = array(i+1,j) - array(i,j)

                if diff lt (array_size+1)/2 and diff gt -(array_size+1)/2 then begin

                        prob_array(diff+(array_size-1)/2) = prob_array(diff+(array_size-1)/2) + 1

                endif

        endfor

; Convert values in prob_array to probabilities and compute entropy

n = total(prob_array)

entrop = 0

for i = 0, array_size-1 do begin

        prob_array(i) = prob_array(i)/n

        ; Base 2 log of x is Ln(x)/Ln(2).  Take Ln of array element
        ; here and divide final sum by Ln(2)

        if prob_array(i) ne 0 then begin

                entrop = entrop - prob_array(i)*alog(prob_array(i))

        endif

endfor

entrop = entrop/alog(2)

return, entrop

end
```

5.5.1.3 COMPRESSION SOFTWARE TESTS

5.5.1.3.1 Description of Tests

5.5.1.3.1.1 Test Objectives

The main questions which the testing of the image compression software sets out to answer are:

1) Do the different PMEM pages perform as expected, i.e. do sky images always compress best under the sky table ?

2) What is the relationship of image entropy to compression ratio and the amount of data loss?

3) How does the GOB length parameter affect compression ratio?

4) What are the advantages/disadvantages of the busy PMEM page with the scaled quantization matrix (PMEM page 1) ?

5) Does the compressor ever perform in unexpected ways?

5.5.1.3.1.2 Test Images

Testing of the compression software was performed using a set of about 150 Galileo images of Venus and the Moon. These images were obtained in their original VICAR format and have had no previous lossy compression or image processing applied to them. Galileo images were chosen because they are CCD images, and hence are as close as possible to what will be obtained with the Cassini cameras. Voyager images, while their targets may be more suitable, were obtained with a Vidicon camera which has poorer performance than a CCD camera. The images were chosen to encompass a wide range of image entropies and image parameters.

An IDL script was written to determine which images qualify as 'busy', 'sky', and 'atmospheric' by checking which PMEM page gave the best compression ratio with a given image (at a B value of 0). Images compressing best under page 0 are 'busy', those compressing best under page 2 are 'sky', and those compressing best under page 3 are 'atmospheric'.

Thumbnail collages of the test images can be found below.

Busy Images:

The following thumbnail images increase in entropy as one `raster-scans' from the top left corner. The entropy ranges from 2.03 to 5.29. The images are all of the Moon.

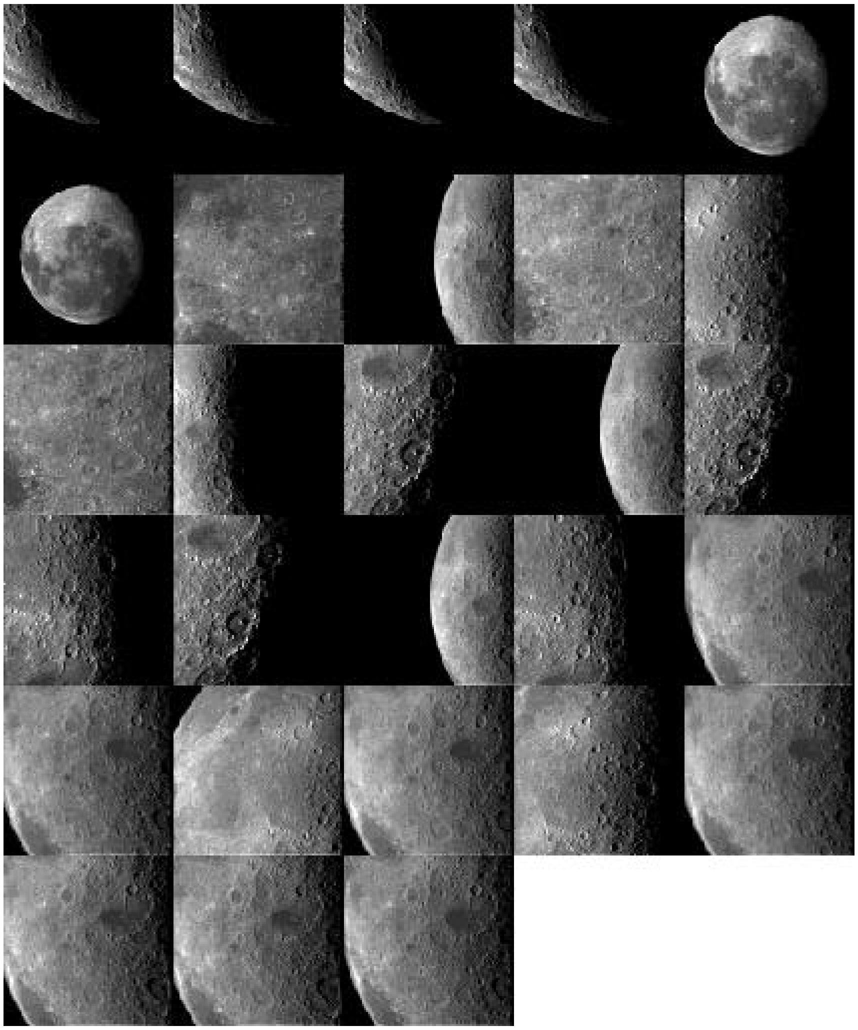

Figure 5.5.1.3.1-1 - Lossy Compression Busy Images

Sky Images:

The following thumbnail images increase in entropy as one `raster-scans' from the top left corner. The second block is a continuation of the first. The entropy of the images ranges from .64 to 2.86. The images are of Venus and the Moon, with some of the Venus images taken at close range.

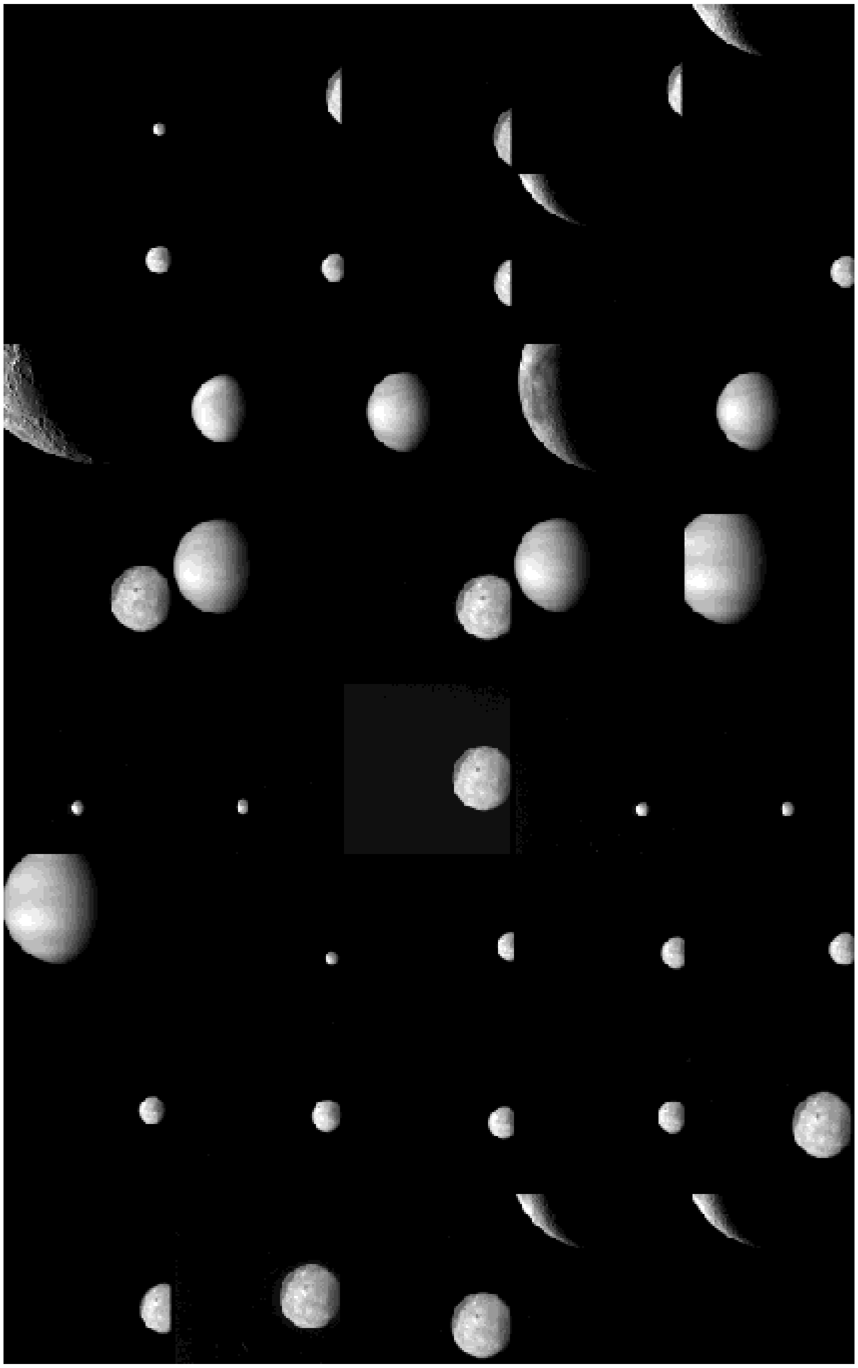

Figure 5.5.1.3.1-2 - Lossy Compression Sky Images

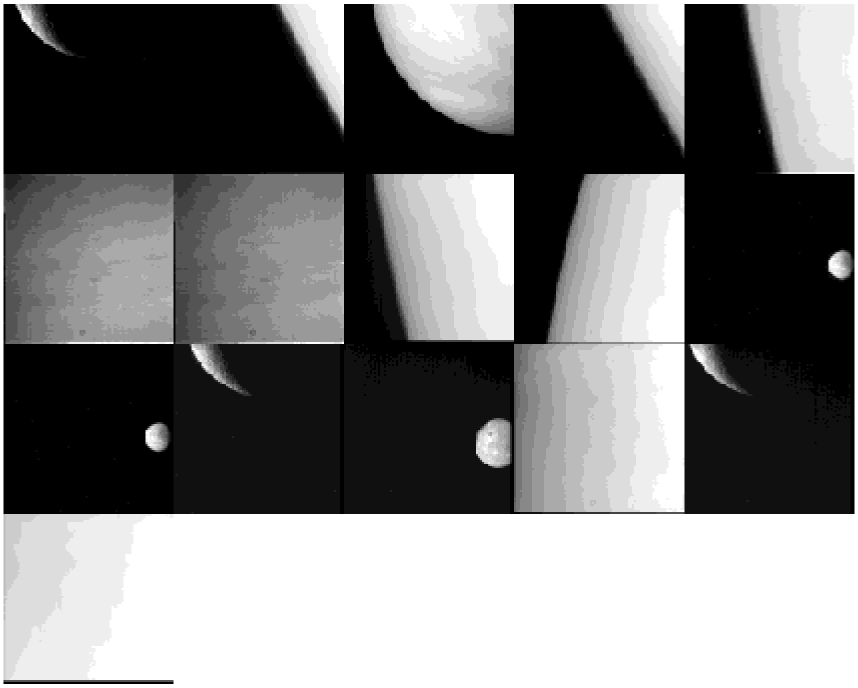

Figure 5.5.1.3.1-3 - Lossy Compression Sky Images (Set #2)

Atmospheric Images:

The following thumbnail images increase in entropy as one `raster-scans' from the top left corner. The second block is a continuation of the first. The entropy ranges from 1.47 to 3.65. The images are of Venus and the Moon, with most of the Moon images taken at low phase angle.

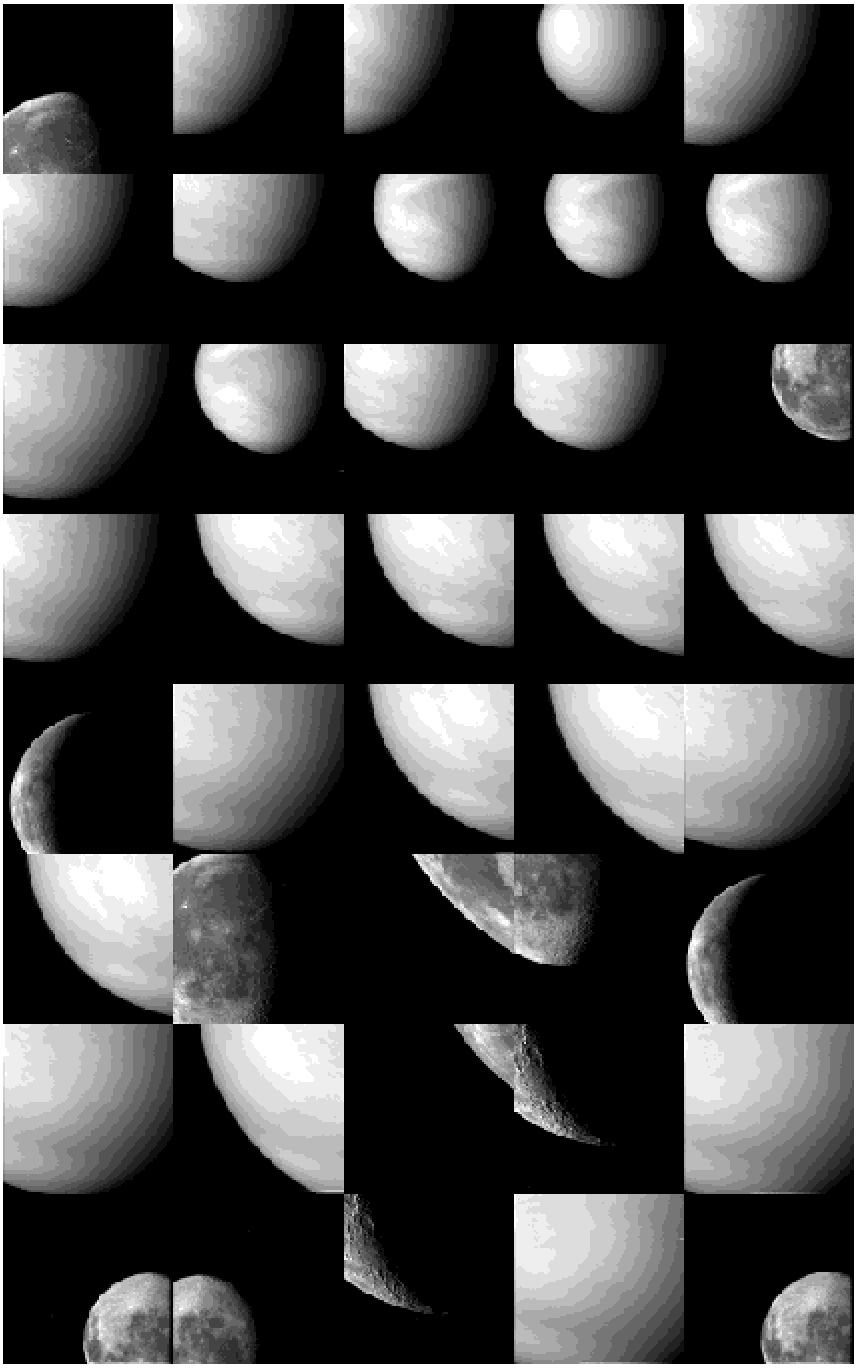

Figure 5.5.1.3.1-4 - Lossy Compression Atmosphere Images

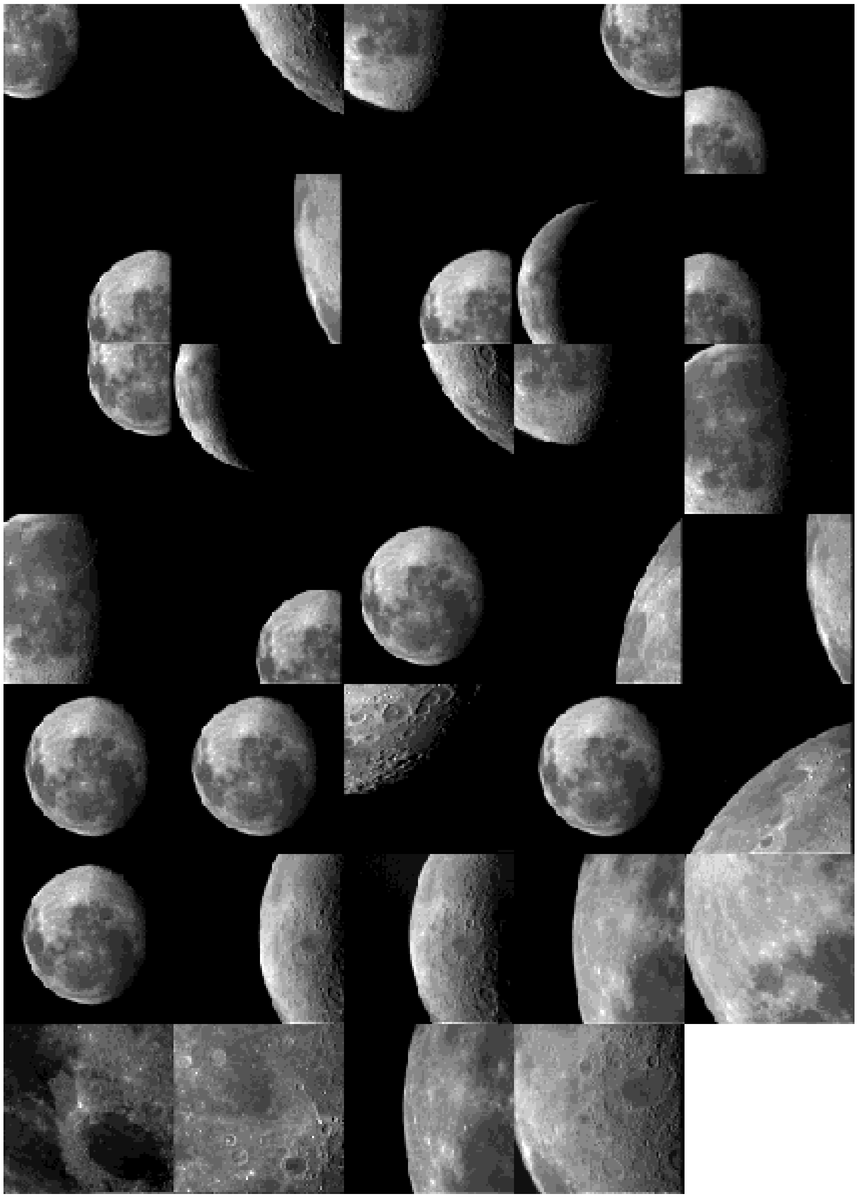

Figure 5.5.1.3.1-5 - Lossy Compression Atmosphere Images (Set #2)

If you are interested in getting copies of the actual raw images I used, please e-mail me at obrien@astrosun.tn.cornell.edu

5.5.1.3.1.3 Test Procedure

An IDL script was written to automate the compression, decompression, and data taking for all of the images. For each group of images (busy, sky and atmospheric), each image was compressed and decompressed with each of the 4 PMEM pages and B values of 0, 1, 2, 3, 5, 10, and 15. A GOB length of zero was used, the NLS bit was set to 0, and optimized decompression was used. The RMS difference between the original image and the decompressed image in each case was found, and is used as a measure of data loss. Also, the compression ratio in each case was read from the compressor log file. Data files were generated for each image type, PMEM page, and B value containing the image names and entropies and the compression ratio and RMS difference value for each image. Once this data was taken, numerous IDL programs were written to graph and analyze the data. The results appear on the following page.

To determine the dependence of compression ratio on GOB length, three image with a range of entropies were chosen and compressed with different GOB lengths. IDL was used to plot the data and determine the relationship. The results of this also appear on the following page.

5.5.1.4 RESULTS OF COMPRESSION TESTS

5.5.1.4.1 Summary:

Since I have collected over 100 data files and could concievably make thousands of plots of the data, this page does not contain every single detail of the compressors performance--there are probably a few details I have yet to stumble upon! This page summarizes the main aspects of the compressor's performance in as concise a way as possible. Reference 5.5.1-1 for data files and analysis software so that, if so inclined, one can examine the data in more detail.

5.5.1.4.2 Do the PMEM Pages Perform As Expected ?

In general, a given image (busy, sky, atmospheric) compresses best under the PMEM page designed for that image type when a B value of 0 is used. This can be seen in the collage of test images (Figure 5.5.1.3.1-1 through Figure 5.5.1.3.1-5). Also apparent in these collages, however, is the fact that some images unexpectedly compress better under a different page than expected. These fall into two main categories:

1) In the set of atmospheric images, there are a large number of pictures of the moon which show up. In general, these images were taken at low phase angle, and hence have a relatively low amount of contrast. Since the atmospheric PMEM page (page 3) is designed for intermediate entropy values (which most atmospheric images fall into), this is not as anomalous as it may first seem. The difference in compression ratio between the atmospheric PMEM page (page 3) and the busy page (page 0) for these images is generally around 2 or less, so these images are borderline cases.

2) In the set of sky images, there are numerous images of Venus in which Venus takes up the majority of the image frame and there is little black sky. There is very little contrast in these images, however, and the atmosphere of Venus appears very smooth (possibly a little out of focus). Since the sky page (page 2) is designed for low entropy images, the fact that these images compress best under the sky page is not surprising. The difference in compression ratio between the sky PMEM page (page 2) and the atmospheric page (page 3) for these images is generally small, so these images are also borderline cases.

5.5.1.4.3 How Does B value Affect the Choice of PMEM Page ?

When a Busy image is compressed at high levels, a large amount of high spatial frequency information is discarded and the resulting data can be more effectively coded for with an atmospheric PMEM page. This effect seems to come into play only at high B values (~10) where the data loss is so high that it is unlikely such values would be used. Even then, the increase in compression ratio is generally small.

A similar effect probably also occurs in which atmospheric images will compress better with the sky page at high B values, but the fact that the sky page has a different set of scale factors than the other pages makes a comparison between the two pages difficult.

5.5.1.4.4 Difference Between PMEM Pages 0 and 1

PMEM page 1, as mentioned before, has a scaled quantization matrix which is designed to discard more high frequency information from an image than the flat matrix in PMEM page 0. While no detailed analysis has been performed thus far, it can be shown that an image compressed under PMEM page 0 can be compressed under PMEM page 1 with a lower B value and have essentially the same compression ratio and RMS difference value. The difference, however, is that compression with PMEM page 1 would better preserve the low spatial frequency components of the image while compression under PMEM page 0 would better preserve the high spatial frequency components of the image. This factor should be taken into account if there are instances when either high or low spatial frequency information is important for certain types of analysis.

5.5.1.4.5 When is an Image a `Sky Image' ?

While no detailed analysis has been performed, a rough estimate can be placed on how much black sky must be present in an image for it to qualify as a 'sky image'.

Moon Images: ~8/9 must be black sky (moon must comprise less than ~1/9 of the image)

Venus Images: ~3/4 must be black sky (Venus must comprise less than ~1/4 of the image)

The fact that Venus can be larger than the moon and still be a sky image is due to the fact that Venus is generally a lower entropy target than the moon.

5.5.1.4.6 Relation Between Image Entropy and Compression Ratio / Data Loss

The general relationship between image entropy and the RMS difference between the original and decompressed image is linear, while the relationship between image entropy and compression ratio is liner in a log-log plot. A program (available on the downloads page) which plots the RMS and compression ratio values vs. entropy and displays the best fit lines and equations. The four plots show the results for each PMEM page and the corresponding image type at a B value of zero.

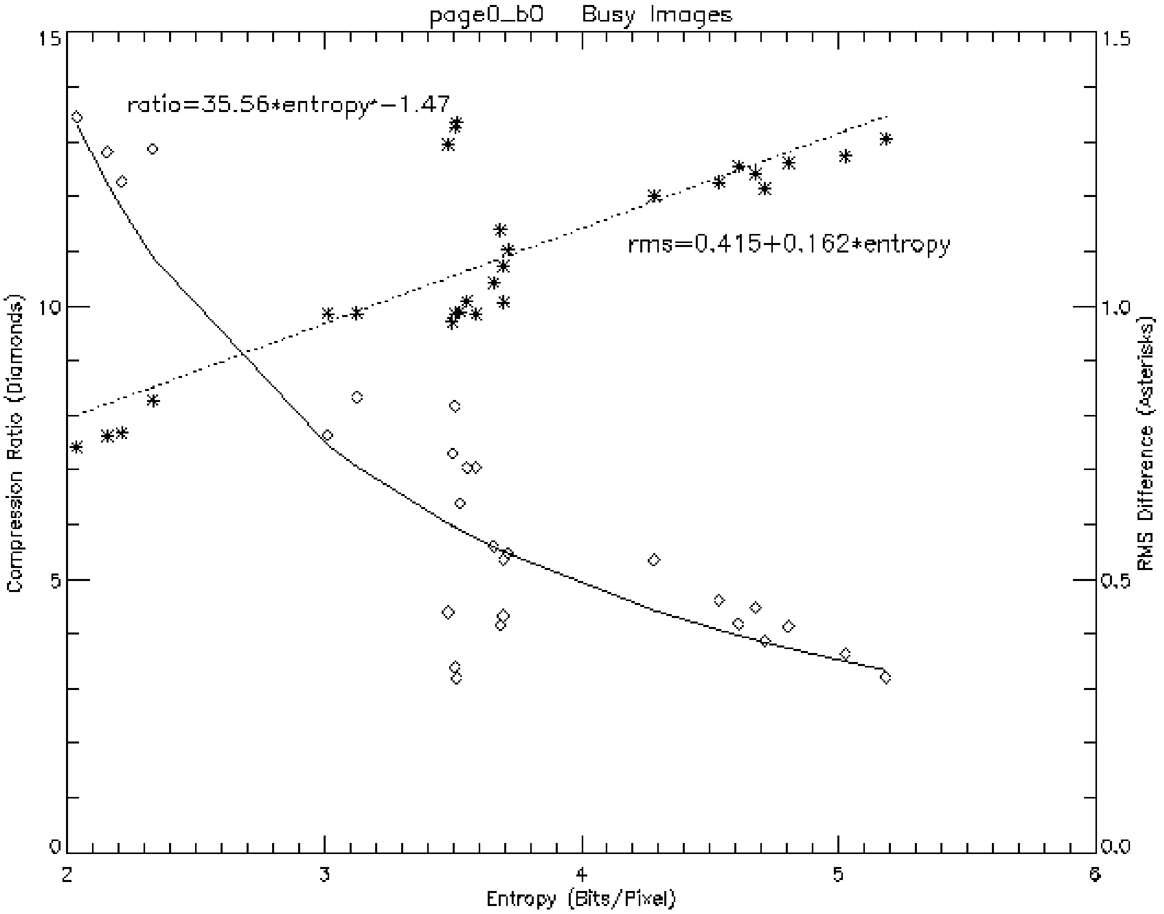

Figure 5.5.1.4.6-1 - Compression Ratio and RMS vs. Entropy (PMEM Page 0, B=0)

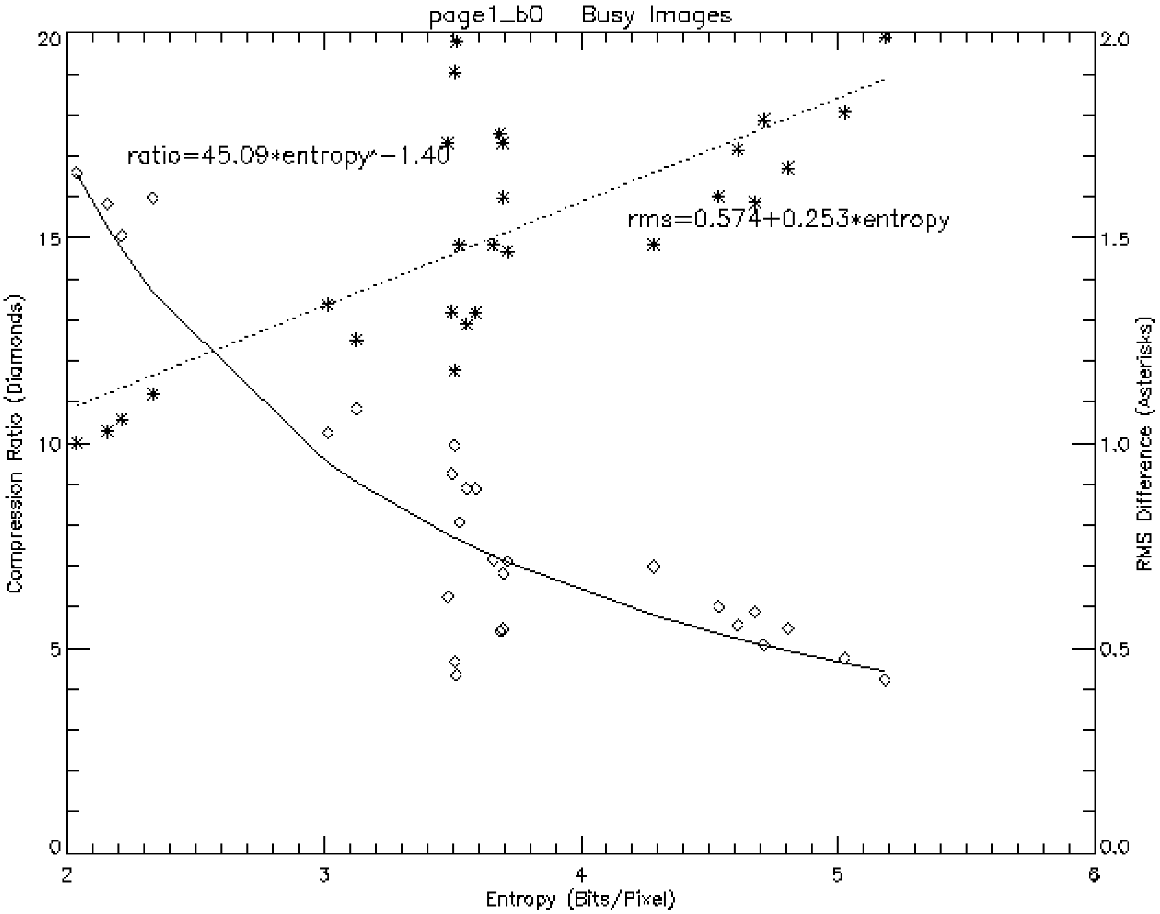

Figure 5.5.1.4.6-2 - Compression Ratio and RMS vs. Entropy (PMEM Page 1, B=0)

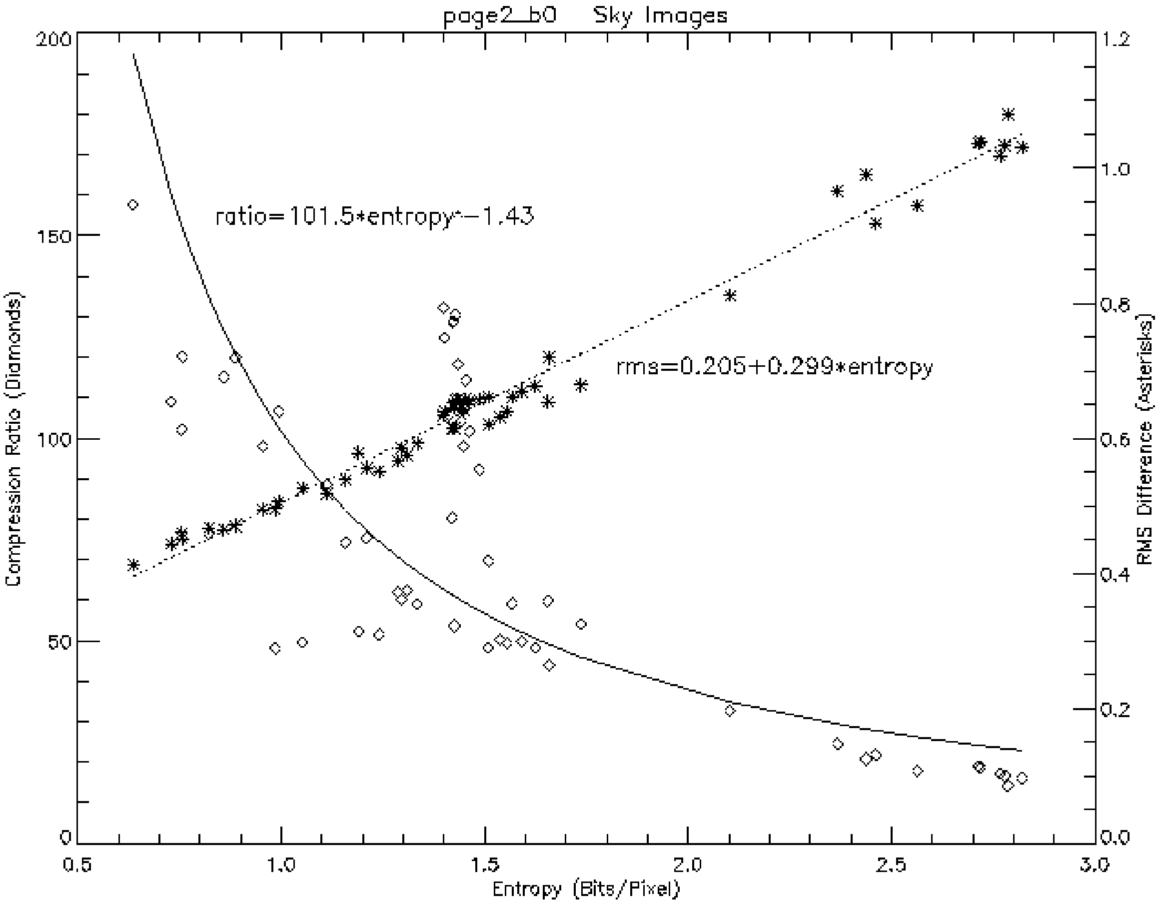

Figure 5.5.1.4.6-3 - Compression Ratio and RMS vs. Entropy (PMEM Page 2, B=0)

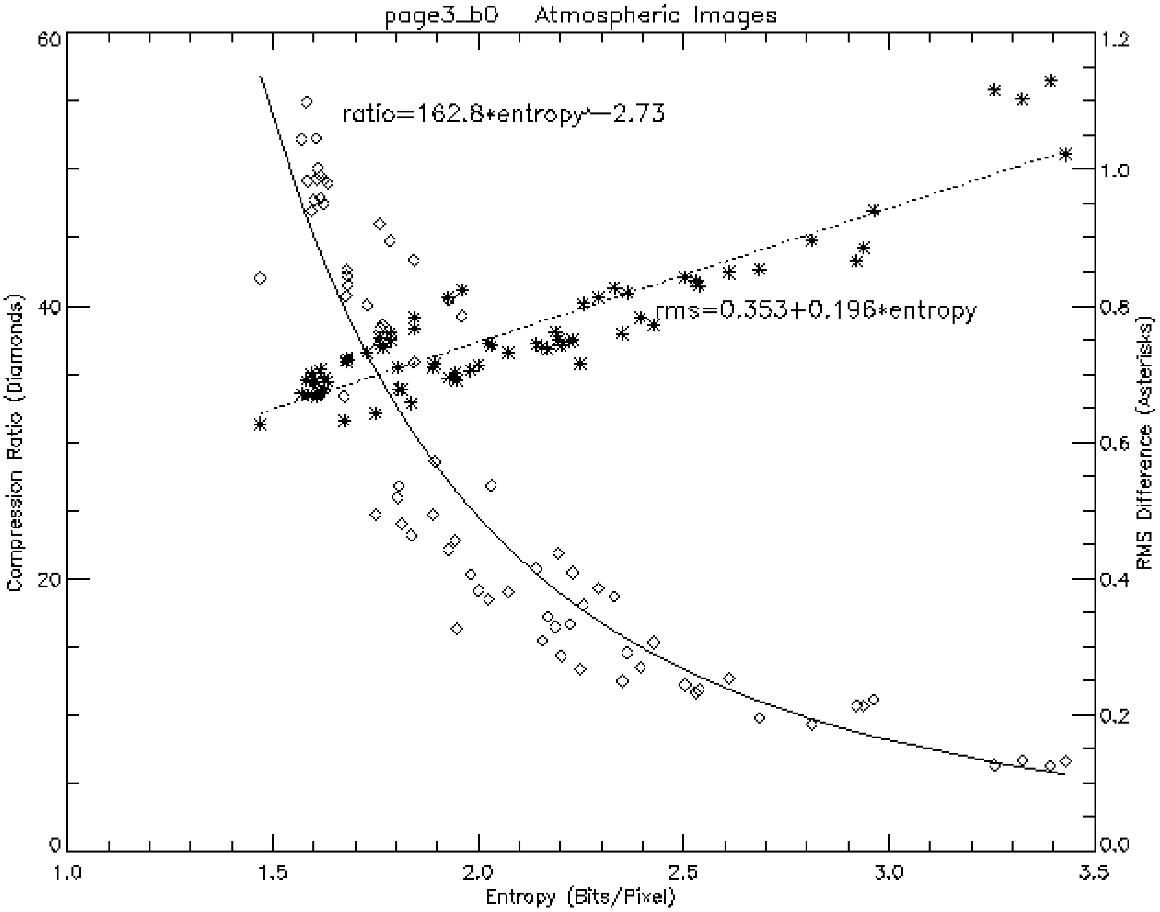

Figure 5.5.1.4.6-4 - Compression Ratio and RMS vs. Entropy (PMEM Page 3, B=0)

There is obviously a large amount of scatter in most of the plots. This indicates that image entropy, as computed here, is not the best predictor of compression ratio and RMS loss. I am currently working on an improved prediction factor which takes into account scene contrast as well as the ratio of black sky in an image. This may improve the correlation and decrease scatter, but it will be impossible to obtain an exact relationship.

5.5.1.4.7 Relation Between B Value and Compression Ratio / Data Loss

In general, the effect of increasing B value on compression performance can be demonstrated with the following plot. It is based on best fit lines to the data for atmospheric images compressed with the atmospheric PMEM page, but plots constructed with data for the other image types have the same general curve shapes.

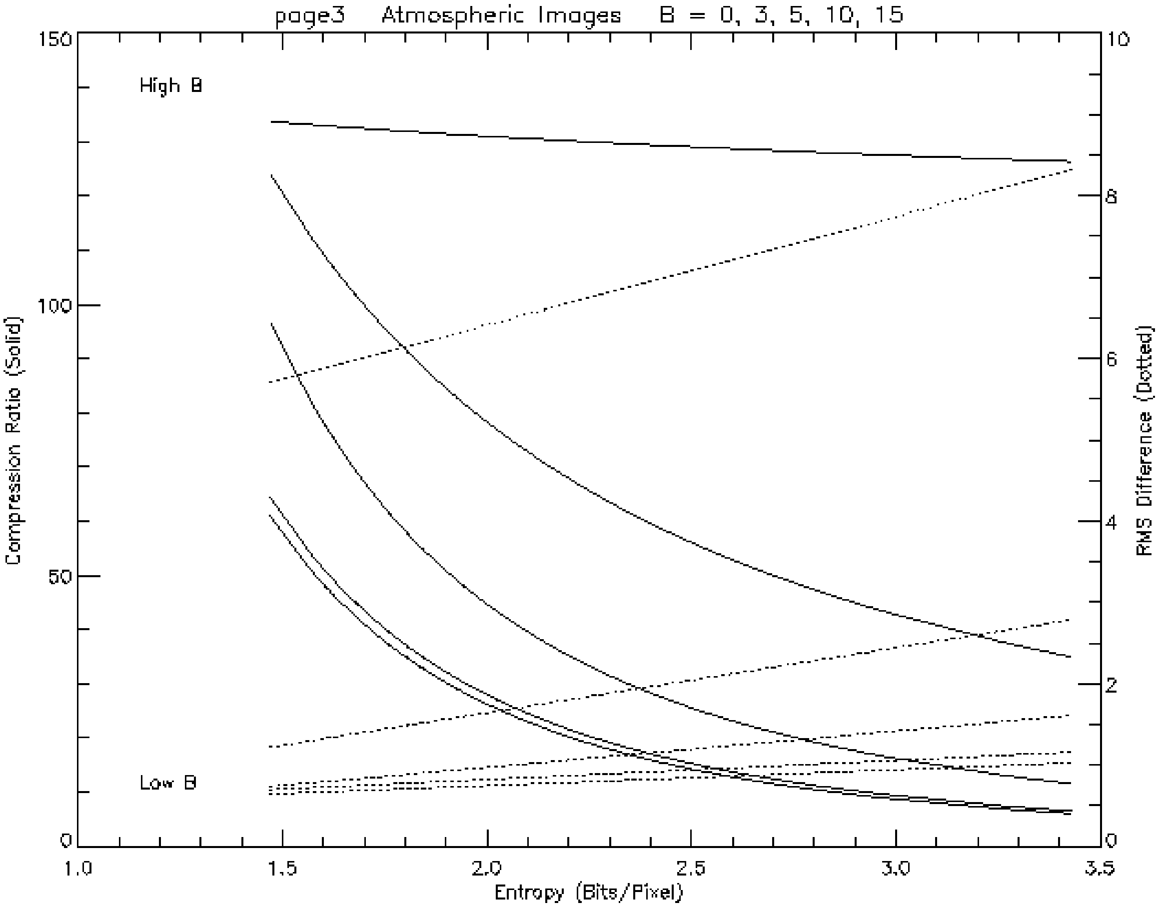

Figure 5.5.1.4.7-1 - Compression Performance vs. B Value (PMEM Page 3)

Generally, the slope of the RMS difference lines increases and the compression ratio curves flatten as the B value increases. Also, the spacing of both the RMS lines and the compression ratio curves generally increases as B value increases. The following groups of plots show the data for specific images from each image type (busy, sky, and atmospheric).

The following plot is for a busy image of the moon compressed under the busy table (PMEM page 0). Essentially all of the busy images I used have the same general compression profiles as in the following plot--a fairly flat response at low B values and a relatively rapid increase in compression ratio and RMS difference at higher B values.

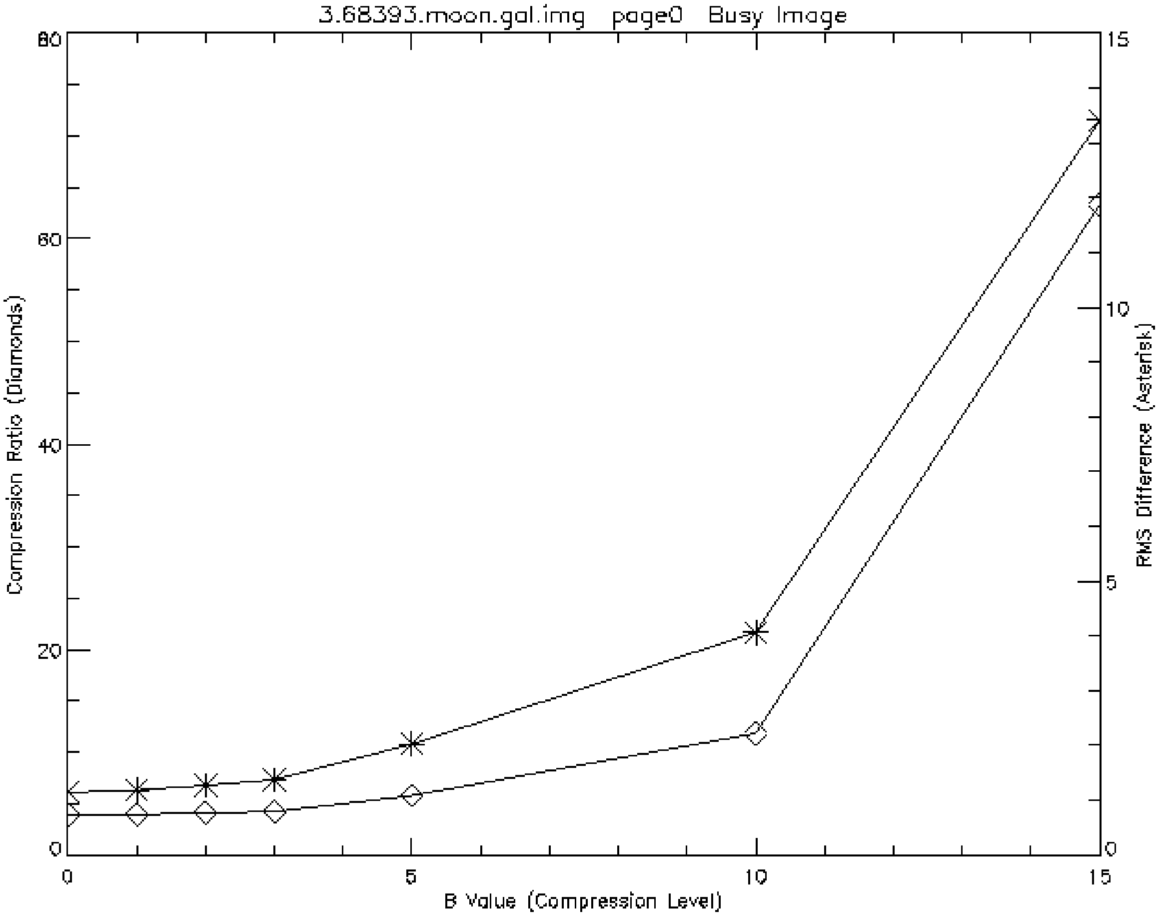

Figure 5.5.1.4.7-2 - Compression Performance vs. B Value (PMEM Page 0)

The following two plots are for a moon and Venus sky image compressed under the sky page (PMEM page 2). In general, sky images have a rather chaotic response to increasing B value. The compression ratio increases with B value, as expected, but the RMS difference value often shows very little increase or an actual decrease as B increases (the decrease is more prevalent in moon images than in Venus images). The decrease in RMS difference does not occur with any other combination of image type and PMEM page. The reason for this is unclear, as increasing B value should discard more information from an image and thus increase the RMS loss. The test procedure and programs were checked multiple times, however, so this behavior is definitely a quality of sky images compressed under PMEM page 2 and not an error in the testing procedure. The implication of this, obviously, is that increasing the B value (from 3 to 5, for example) in the case of sky images will almost always give better compression performance (higher ratio with the same or lower amounts of loss).

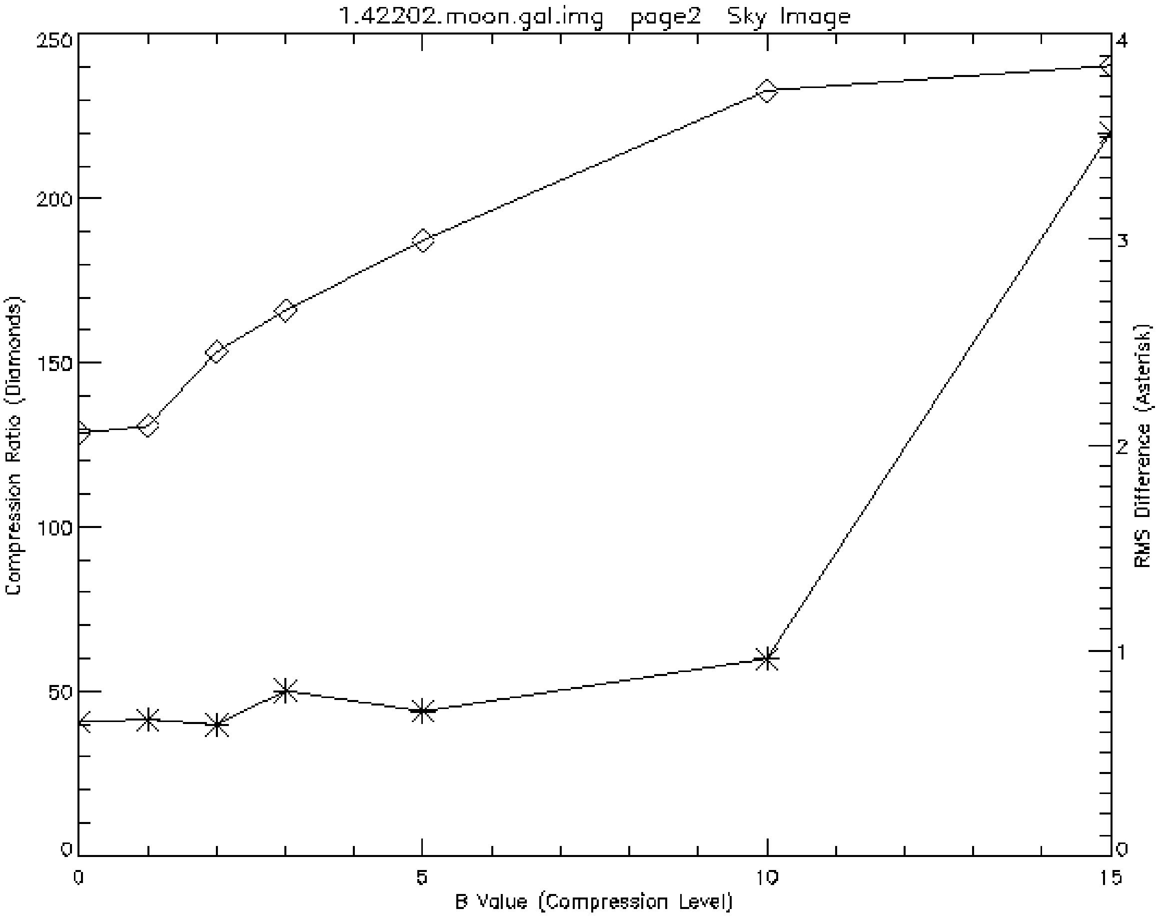

Figure 5.5.1.4.7-3 - Compression Performance vs. B Value (Moon, PMEM Page 2)

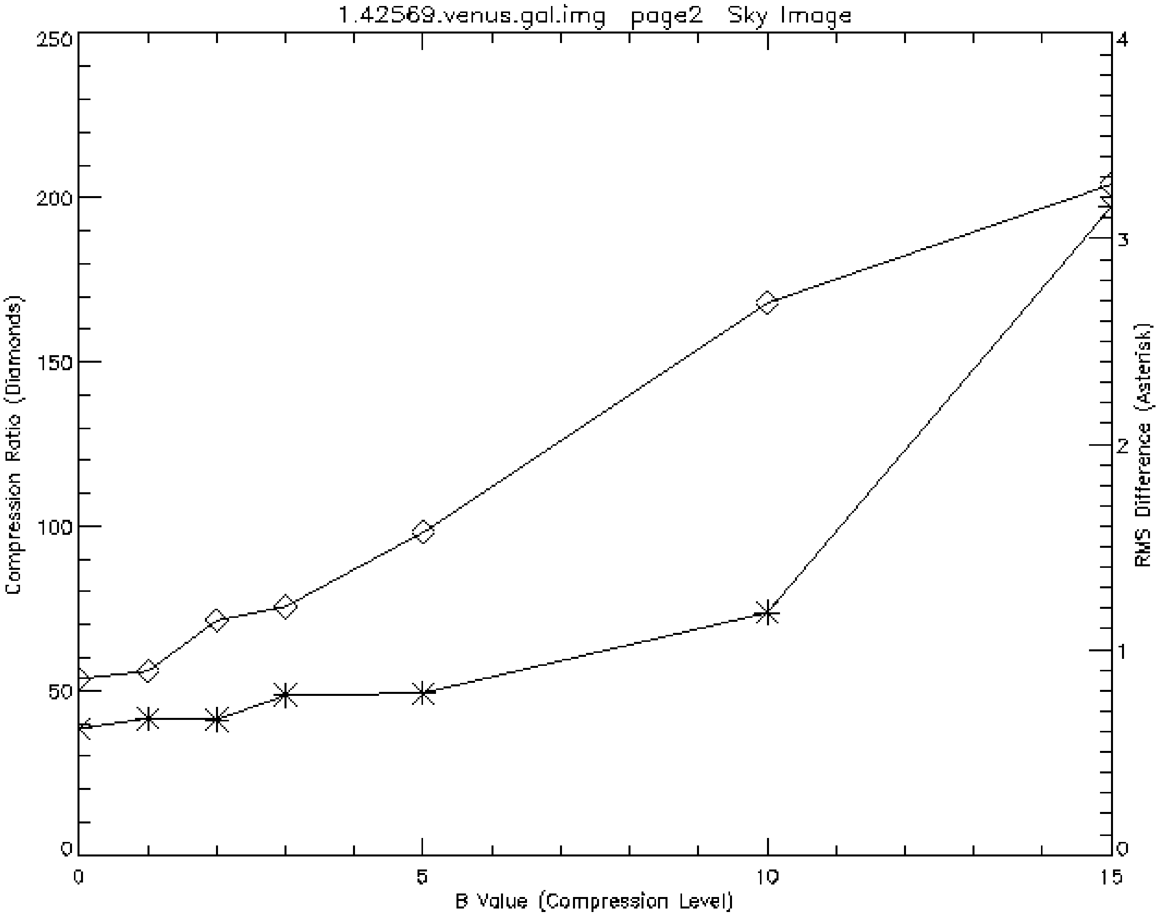

Figure 5.5.1.4.7-4 - Compression Performance vs. B Value (Venus, PMEM Page 2)

The final 2 plots are for a Venus and a moon atmospheric image compressed under PMEM page 3 (the atmospheric page). At low phase angle, remember, the moon compresses best under the atmospheric PMEM page. For moon images, the ratio and RMS curves are very similar to the curves for busy images. For Venus images, however, there is generally a sharp jump in compression ratio between B = 3 and B = 5 and the RMS curve stays relatively flat up to a B value of 10. This sharp jump in compression ratio occurs with only a small increase in RMS difference, so it would generally be advantageous to use a B value of 5 rather than a B value of 3 for imaging targets similar to Venus (Saturn and Titan?).

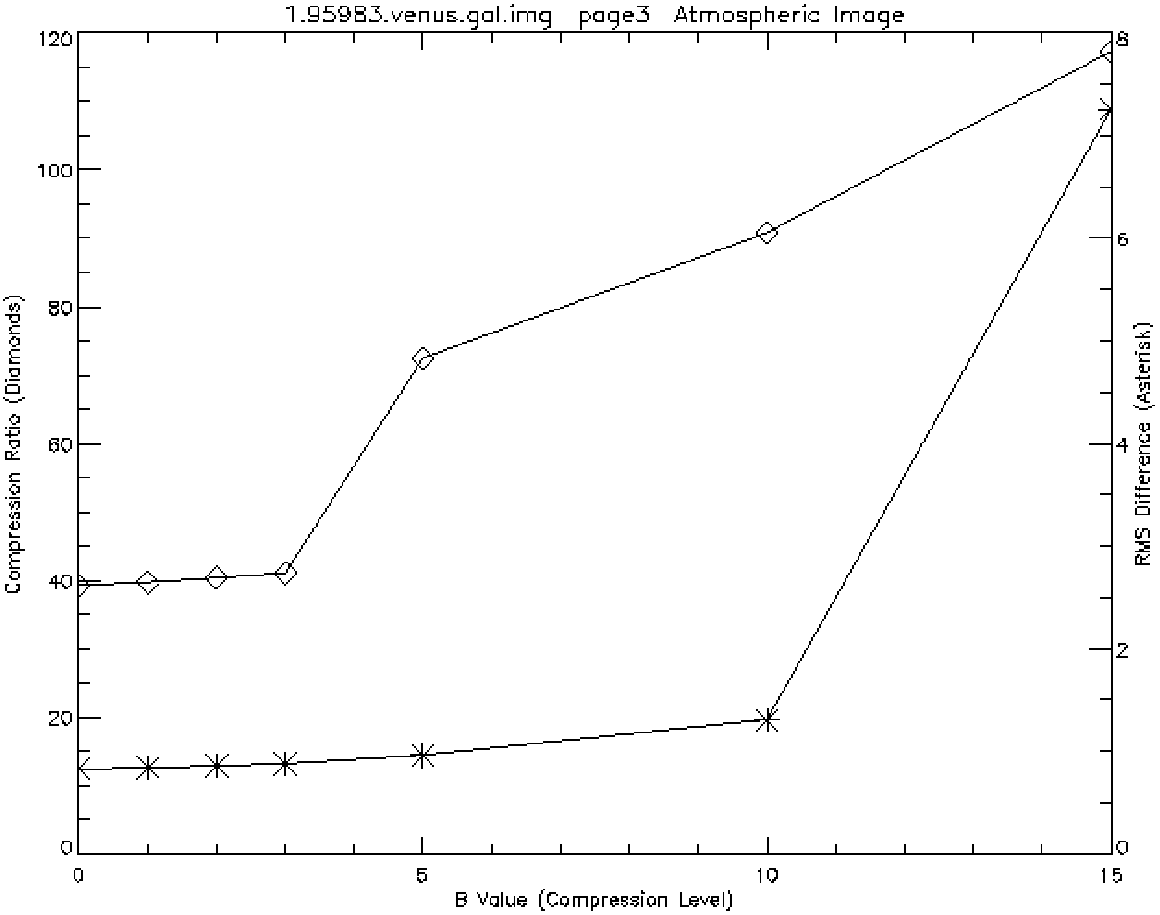

Figure 5.5.1.4.7-5 - Compression Performance vs. B Value (Venus, PMEM Page 3)

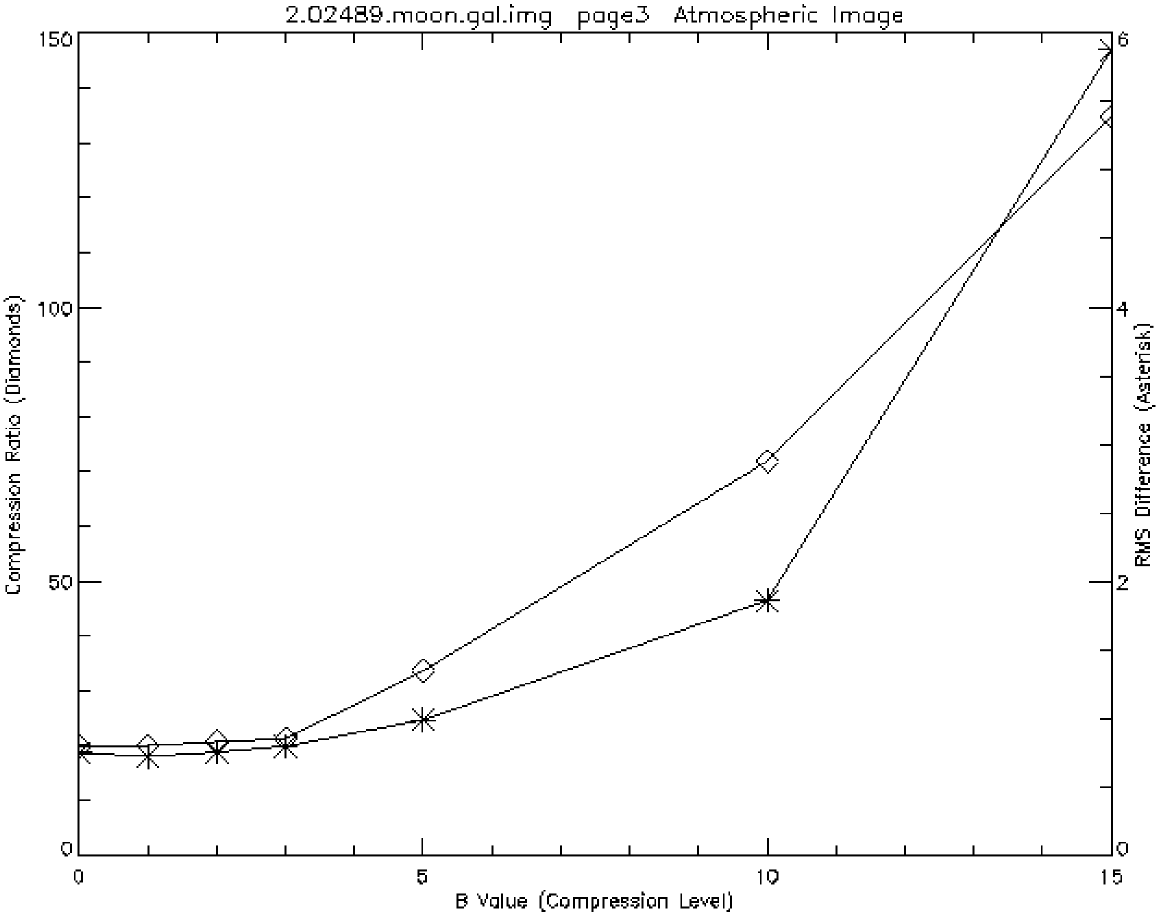

Figure 5.5.1.4.7-6 - Compression Performance vs. B Value (Moon, PMEM Page 3)

5.5.1.4.8 Effect of GOB Length on Compression Ratio

As the GOB length decreases, more error correction headers are included in the compressed image file and the total compression ratio decreases. The plot below shows the actual data for three images of different entropies compressed with a range of GOB values. Once the GOB length drops below about 50, there is a sharp dropoff in compression ratio.

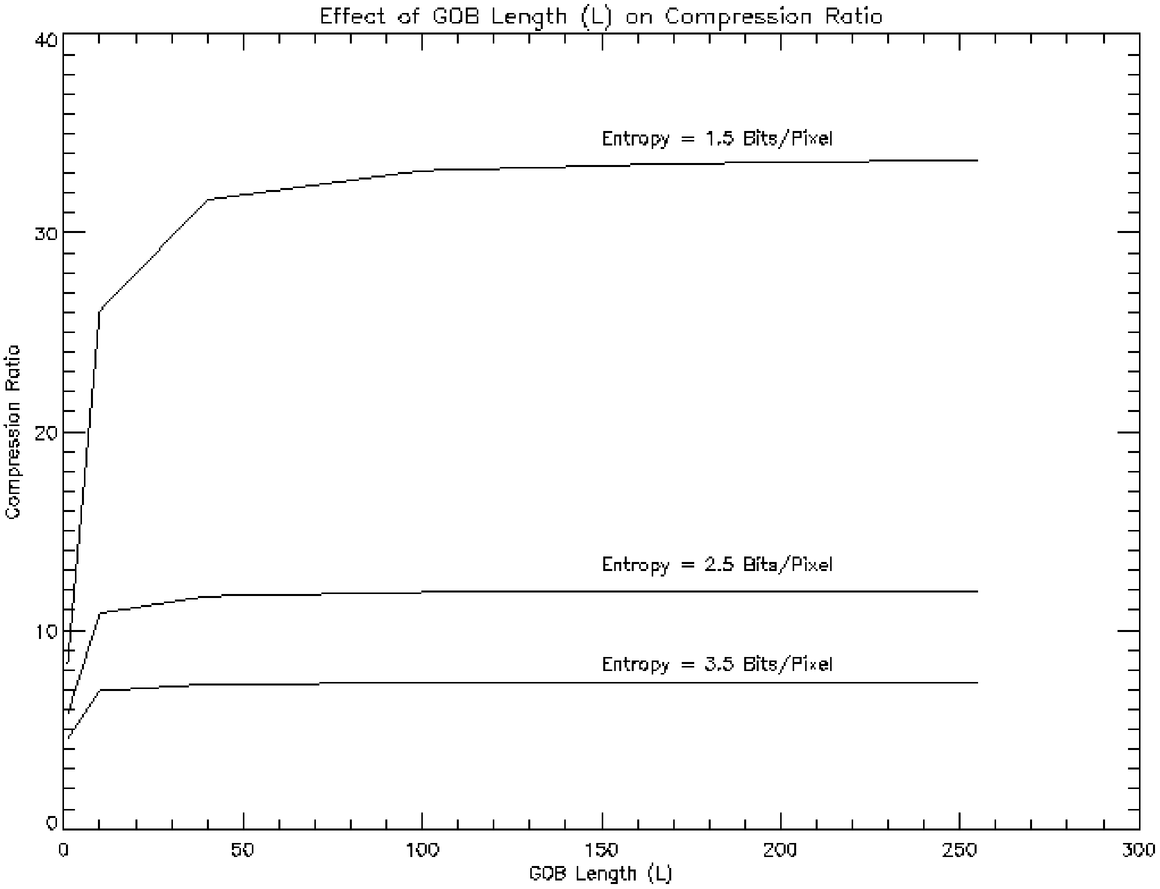

Figure 5.5.1.4.8-1 - Compression Ratio vs. GOB Length

The next plot of the same data is done with the reciprocal of the compression ratio plotted against the reciprocal of the GOB length (L). A linear relationship exists when the data is plotted this way. The Y intercept is the reciprocal of the compression ratio with no GOB headers. The slope is a constant ranging from .080 to .095. I have thus far been unable to find a direct relationship between the slope and any of the other factors involved in compression.

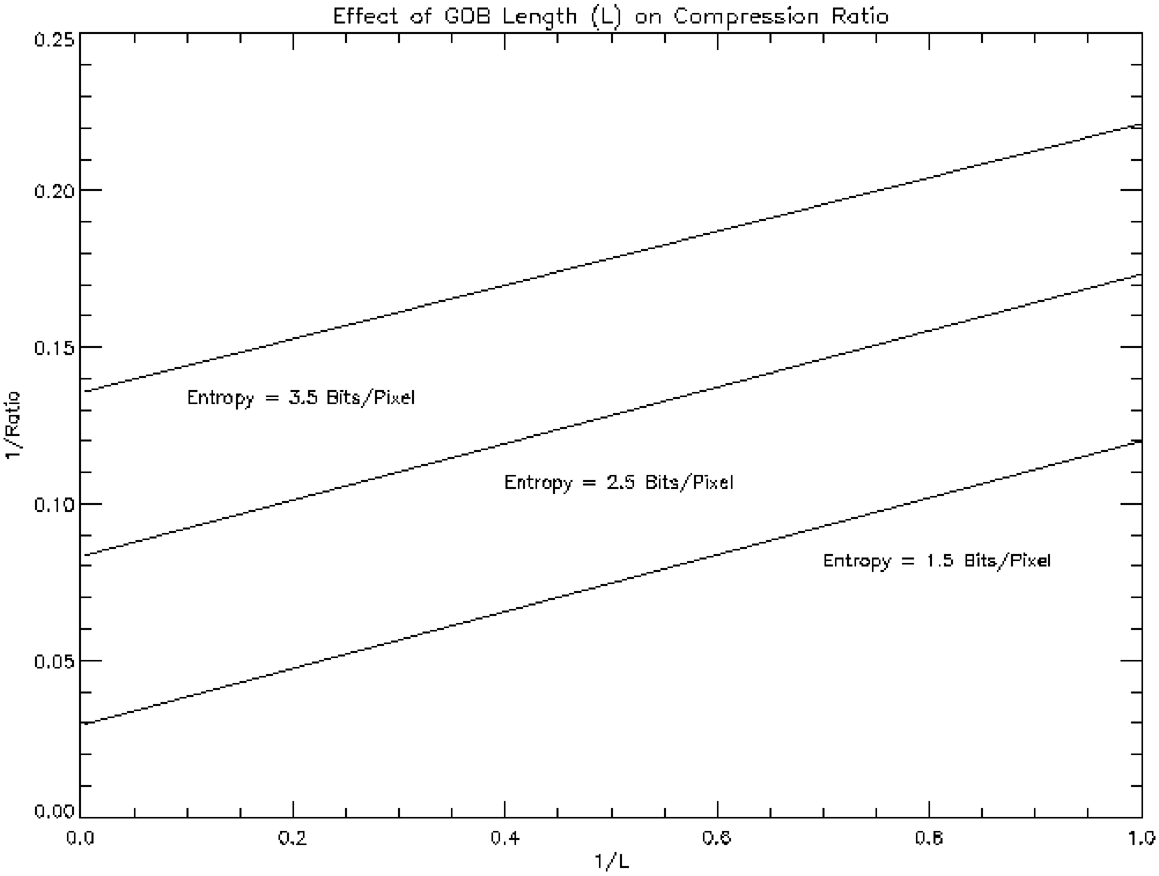

Figure 5.5.1.4.8-2 - (Compression Ratio)-1 vs. (GOB Length)**-1**