Effects of Data Compression on Cloud Motion Measurement in Planetary Images

A REPORT

from the space science and engineering center
the university of wisconsin-madison
madison, wisconsin

THE SCHWERDTFEGER LIBRARY
1225 W. Dayton Street
Madison, WI 53706

Effects of Data Compression on Cloud Motion Measurement in Planetary Images

by

Robert Krauss
Space Science and Engineering Center
University of Wisconsin - Madison

Abstract

This study examines the effects of image data compression on the visual recognition of planetary cloud features, and the quality of cloud motions measured in compressed images. In cases where the compression factor is 6:1 or less, relatively minor degradation was observed in human ability to recognize and interpret cloud features, and identify regions of cloud motion. The ease of computer cross-correlation to obtain accurate measurements of such motions was also affected in minor ways. Typical additional random errors introduced by compression of less than 6:1 were on the order of <1/2 pixel, or <2 m/s rms for the resolution in the tested image data. There was also a 2-3% decrease in the number of successful correlations. At the higher 10:1 and 16:1 compression factors, some visual changes in the images were apparent, and the random error in position measurement was increased by 1.0-1.5 pixels (corresponding to ~3-5 m/s added rms dispersion in velocity). At these higher compression factors, the probability of misinterpreting some successful correlations as good motion vectors, when the targets failed to correlate successfully in uncompressed images (and thus should be rejected as bad), also became significant. Compression factors above 10:1 should be used with caution, and preferably on familiar data where the quantitative effects of compression can be reliably predicted.

[1981]
[July or August]

I. INTRODUCTION

This study was undertaken at the Space Science and Engineering Center (SSEC) under contract to Jet Propulsion Laboratory to investigate the effects of using data compression on planetary cloud images. Compression of the image data on future spacecraft could be a very useful means of increasing the number of images returned when low power and large distances limit the communications bandwidth. To be useful for atmospheric circulation studies, however, the images can only be degraded by compression in a way which still preserves the essential characteristics of the clouds and their motions. The question which must be answered here is how seriously the integrity of the image data is compromised with respect to (a), the visual recognition of cloud features, and (b), the quality of measurements of cloud motions.

II. DATA USED

Two sets of images were used for this analysis. The first set consisted of five Mariner 10 UV images of Venus. The second set consisted of four each of orange and green filter Voyager 1 images of Jupiter. Orange and violet were originally specified to provide maximum color contrast for measuring winds on Jupiter, but green images were received at SSEC instead of the violet, and since green and orange Voyager Jupiter images are so much alike in effective filter wavelength and cloud texture, no color dependent results can be presented here. Indeed, the orange Jupiter images received very little analysis.

Most image data is somewhat redundant, having more bits per pixel than are necessary to totally and unambiguously describe the information content. It is possible to take most images with 8 bits/pixel, reduce the average bits/pixel to a value somewhere between 2 and 6 by a coding scheme, and then decode the image back to an exact copy of the original. The reason for this is that the DN value of any pixel is likely to be close to the DN value of its neighbors, since an image generally contains objects larger than a pixel in size. One can define a quantity called "data activity" or "entropy" (measured in bits/pixel) as a means of quantitatively describing the extent of such correlation between neighboring pixels. A low mean entropy for an image implies a high degree of correlation, while high mean entropy implies more randomness in the pixel values. Venus and Jupiter images range from 2.8 to 3.1 bits/pixel.

Two different coding or compression schemes were used in this study. For the compression factor of 2.5:1 a Block Adaptive Rate Controlled (BARC) compressor was used. An activity accumulator works while the image is being loaded to buffer. This makes it possible to select a compression factor based on the measured image entropy, or alternatively, a compression factor arbitrarily set by the user. Based on the number of allowed bits per line, the pixels are coded at the appropriate factor.

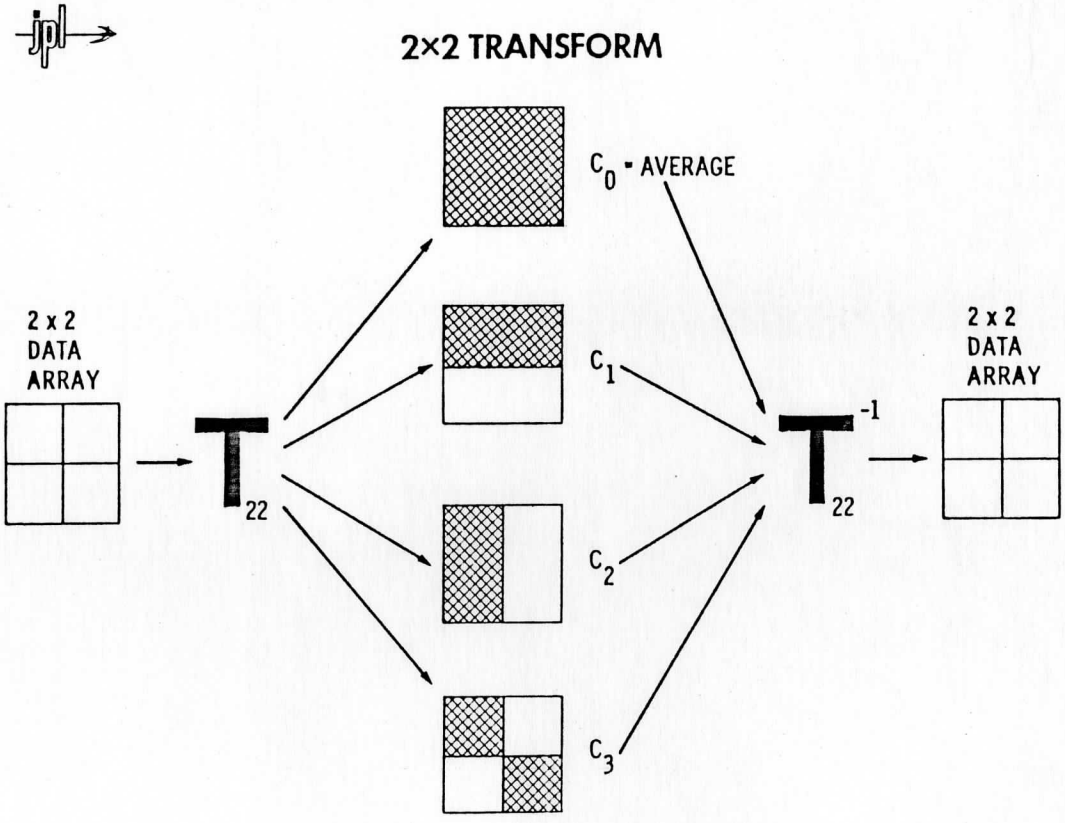

Figure 1. Schematic diagram showing the lowest order RM2 Hadamard transform decomposition of a sub-image domain into averages representing different two-dimensional spatial frequency components. These averages are edited and coded to comply with the imposed bits/image limit prior to transmission from spacecraft to earth by means of a highly noise free transmission link. The inverse transform, done after the image is received on earth, decodes and reconverts the pixel averages back to a close representation of the original image.

Since the BARC compression is not two-dimensional, it doesn't work well at the higher compression factors used in this study. In those cases (4:1, 6:1, 10:1, 16:1) an RM2 subpicture compressor was used. The RM2 compressor, consisting basically of cascaded Hadamard transforms, works on a 64 x 64 pixel subimage, calculating pixel averages in square sub-domains. Based on the allowed number of bits per image, the RM2 averages are edited and subsequently coded at the appropriate factor [SPIE Proceedings, 66 ,pp 70-89 (1975)]. At rates less than 0.8 bits/pixel, a 2:1 edit is performed prior to decomposition into local averages. This would have an effect similar to reducing the image resolution by a factor of 2 in each direction. By using primarily averages and edits, the compression algorithms can be kept simple enough to operate very fast, and be hardwired on the spacecraft.

All images were compressed and decompressed at JPL with compression ratios ranging from data preserving to 16:1, where obvious visual differences could be detected. Prior to compression, the images were "cleaned up" by removing noise spikes, and in the case of the Voyager images, locating and removing the reseau marks. The decompressed images were photometrically corrected, after which a uniform contrast stretch was applied in each case. The images were then geometrically rectified to remove inherent distortion in the vidicon camera, using the reseau locations determined prior to the image compression. Except for the compression and decompression, the images used as input to this study at SSEC were processed exactly as they would normally be processed prior to our standard cloud motion measurement procedures.

The image tapes, as received from the Image Processing Laboratory at JPL, were reformatted on a Univac 1100 at the Madison Academic Computer Center and then input to the Man-computer Interactive Data Access System (McIDAS) at SSEC. Final image processing and video display was done using McIDAS. McIDAS Processing is done interactively, with the image data stored on digital disk at the computer. The operator uses an interactive terminal capable of RGB video and graphics display, with a joystick operated cursor and keyboard/CRT for operator feedback and processing control. Interaction between the image data and operator is via high bandwidth video channels, while interaction between computer and the image data is kept in the precise digital domain for maximum accuracy. Thus the operator guides the computer through a series of complicated quantitative processes, making decisions at critical points as needed.

The RGB video signal at the terminal can be tape recorded on cassette after conversion to an NTSC composite video signal. The analog video display is thus preserved in a form very similar to that seen by the terminal operator during the analysis. This is useful for presenting the visual comparisons and time lapse studies here, and a videotape showing portions of the image comparison analysis forms part of this report.

III METHOD OF ANALYSIS

Five procedures were used in the evalation of the images, covering a wide range of processes of varying degrees of objectivity which might be applied to the images at some point in their scientific analysis: 1. Visual Comparison, 2. Histogram Analysis, 3. Image Differencing, 4. Time Lapse Display, 5. Cloud Motion Measurement. These evaluation procedures are described in the following paragraphs.

1. Visual Comparison Method.

A visual comparison of various compressed versions of the same image can best be accomplished by blinking back and forth between the compressed and uncompressed version of any single image. Errors introduced by the compression are perceived as motions or as variations in brightness. By using the same image in both cases, the true motion of the cloud is suppressed, and only motions introduced by the data compression are seen. These motions represent changes in the position of cloud features. Any change from the true positions is bound to have an influence on the measurement of motion. If the displacements introduced by compression are truly random, the influence on cloud motion measurement will be merely the introduction of a larger dispersion.

2. Histogram Analysis Method.

Gross variations in brightness or contrast in the image as a whole were examined using a histogram of the number of pixels at each DN brightness level. If the compression algorithm alters contrast or brightness to a large extent, then the shape, amplitude, or width of the histogram would be modified.

Small scale variations in brightness can be examined by looking at the cloud features selected by the McIDAS operator as targets to follow for wind field measurements. One would hope to see little, if any, change in the mean DN or brightness of a cloud target or in the rms deviation from the mean (local contrast) for that target as a function of data compression factor.

3. Image Differencing Method.

Data compression changes the original image by either adding or subtracting from the value of each pixel. The best way of seeing where changes occur as a result of data compression is to subtract the original image from the compressed version. Since the differences are generally in the 0-6 DN range for the compression factors considered in this study, it is necessary to boost the contrast of the difference image 20-30 times prior to display. The clouds should be completely removed by the subtraction, leaving nothing but a

flat gray field. This would be expected to occur to some extent in parts of pictures subjected to low compression factors (2-6). Higher compression factors show a noise pattern random in character, but somewhat correlated with neighboring pixels, since the compression algorithm is dependent on spatial frequencies existing at each region of the image. The noise pattern in compressed images is not like the randomly positioned noise spikes seen in images having random bit errors in the data chain. Motions observable in the visual comparison test (test #1 above) can be correlated with localized regions of coherent "compression noise" seen in the differenced images.

4. Time Lapse Display Method.

If the noise pattern in the differenced images is not completely random, a time lapse loop of the differenced images will show true cloud features moving. Random noise added to images would show no pattern or recognizable cloud details. The blink comparison described above shows position changes introduced by compression. Time lapse imagery shows if there is any motion produced by these position changes, i.e. if the position changes are not truly random but instead biased in one particular direction to generate a false motion. Additionally, if we see a low contrast duplicate of the original image or outlines of features in the differenced images, this is evidence that information is being removed in the compression process, and consequently that the cloud motion measurements are being degraded in some way.

5. Cloud Motion Measurement Method.

The fact that compression noise or cloud motion can be added and image information removed from images requires that we measure the motions of identical cloud targets in both compressed and uncompressed images and compare the results to determine quantitatively the effects of compression on the measurement process. Cloud motions are generally measured by the objective process of computer cross-correlation, which can achieve sub-pixel accuracy. It is likely that cross-correlation will be affected in several ways:

a) Cloud motion measurement will be made more difficult by loss of image detail due to compression.

b) The result of a successful correlation may be less accurate in a compressed image due to introduction of compression noise that changes cloud positions.

c) Spurious motions may crop up as a result of compression noise and be interpreted as good velocity measurements when, indeed, they are not.

The ratio of successful measurements to measurements attempted is a way of quantitatively estimating the difficulty of correlation. To see how the accuracy of individual successful measurements is altered as the compression ratio increases, we look first to see if the gross features of the atmospheric circulation are preserved. Then we examine if the rms scatter or dispersion of the wind velocities measured in the compressed images gets larger with higher compression factor relative to the wind velocities measured for the identical cloud feature in the uncompressed images. Motion vectors can also be individually and collectively examined for biases with regard to location, to cloud morphology, and to brightness and texture of the cloud targets. Because of the complexity of the wind motion measurement analysis, that portion will be treated in a separate section of this report.

IV. RESULTS OF IMAGE COMPARISON ANALYSIS

The Venus images, having somewhat less contrast and more diffuse cloud features might be expected to be a better candidate for data compression. This was found to be the case as far as visual comparison is concerned. (Cross-correlation wind measurement biases, as we shall later see, appear to be slightly worse for low contrast features.) Even at a compression factor of 16, the visual differences were subtle and required careful searching and contrast stretching. One had to actively search for differences because they were not visible at a casual glance. Jupiter has more contrast and a much greater variety of cloud textures than Venus. Compression errors could more easily be seen to cluster in regions of high contrast gradients and higher spatial frequencies, regardless of the compression factor. Higher compression factors merely introduced compression noise over a larger area and increased the amplitude of compression noise which had appeared earlier at lower compression factors.

1. Visual Comparison Results.

Both Venus and Jupiter showed extraneous cloud motions and some local shifting of brightness contours as a result of data compression. The biasing of the data was barely discernable at 4-6 compression factors, more discernable at a factor of 10, and very easy to see when a compression factor of 16 was used. Some linear features began to break up, and contouring became noticeable on a smaller scale in low contrast regions at the 16:1 compression factor. Clearly, at higher compression factors (10-16) it would be more difficult to select cloud targets and follow them in time, but it would not be 10 to 16 times harder. Subjectively, we would estimate a drop in the number of cloud targets which could be found and tracked and there would be a larger dispersion in the ability of the terminal operator to define

their position in subsequent images. At 16:1 compression, we estimate that only 10-30% fewer targets would be tracked and the positional accuracy would degrade by a pixel or perhaps two. For determining the gross characteristics of the scene, there would be little reason not to use a high compression factor.

We did not perform any tests to determine the quantitative effects of data compression on a person's ability to select or track clouds. It is important to discriminate between a person's perception of what a cloud is and where it moves, and the computer's "perception" of the same event via cross-correlation. The two are similar, but differ in subtle ways which are not understood.

2. Histogram Analysis Results.

The gross shape, height, and width of histograms of brightness distributions were preserved even at the highest compression ratios, with only minor differences in degradation observed. Generally speaking, the changes were that neighboring bins might vary a few percent in their contents, but that any group of contiguous bins tended to preserve its total. The bin content deviations increased as the compression factor increased, but never got higher than 5% in any one bin, even for a compression factor of 16. Data compression must obviously be a relatively local phenomenon with no major effects on large scale image contrasts and brightness. This result is expected, given the averaging technique used in RM2 compression.

We then looked at small samples of data in 15 x 15 pixel regions to see if there were smaller scale differences. The mean DN and rms deviation from the mean were calculated. The differences of the mean DN and rms deviations from the mean DN between the various compression factors and also between the compressed and uncompressed images were of the order of a few tenths of a DN. The rms deviation tended to decrease slightly at higher compression factors while the mean DN remained almost fixed. We conclude that, locally at the cloud target scale, data compression has very little quantitative effect on the data. Some of the image details may be lost.

The result of the histogram analysis indicates that in 15 x 15 pixel regions, or larger, data compression at factors 16:1 or less has no significant effect on the ability to do photometric studies or characterize cloud targets by means of average brightness or color ratios. We did not examine regions smaller than 15 x 15 pixels.

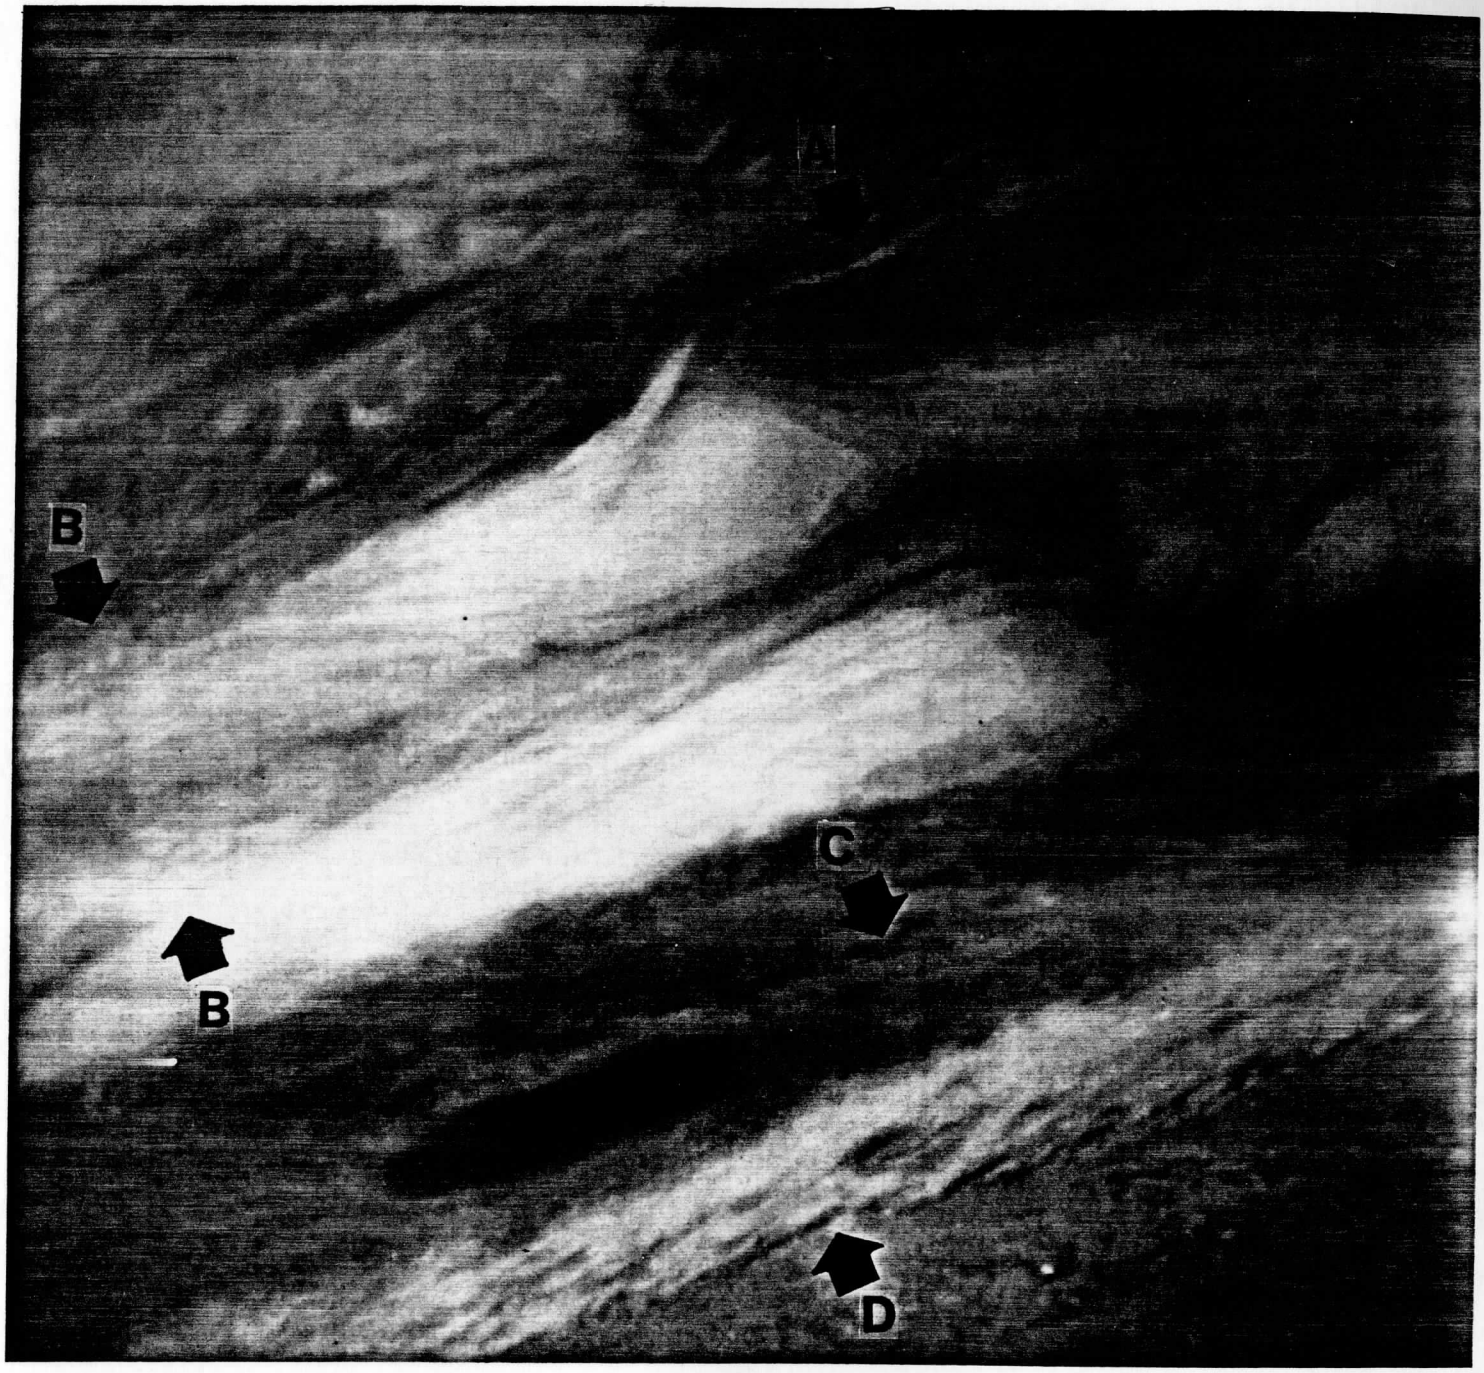

Figure 2a. Portion of one of the uncompressed images of Jupiter used in this study. Note the features indicated by the arrows, and compare with Figure 2b and Figure 3d (which represents the difference of 2a and 2b). The effect of data compression is most easily seen in the changes of linear features, which tend to be broken up at higher compression factors.

Figure 2b. Portion of 16:1 compressed image. Comparison with Figure 2a shows some features such as (A) which are broken up into segments. Features at (B) show more contouring and less intricate small details. The same is true at (C). The high spatial frequency details at (D) are less distinct. Comparison with Figure 3d shows that many of the high frequency details have been lost in the compression process.

3. Image Differencing Results.

Two major characteristics of the differenced images (see Figures 2-3) were observed:

1) The "compression noise" was organized into 72 x 72 pixel boxes with discernable borders, which we took to be an artifact of the RM2 compression process, and

2) the errors tended to cluster in groups which grew in size and intensity as the compression factor increased. The second characteristic is much like that of bacteria growing in a dish around a food supply. These characteristics would seem to support the view that there are certain properties of the data that the compression algorithm has trouble handling, and that the compression errors are spatially correlated with one another to some extent. The clusterings generally occurred around regions in the images characterized by high brightness gradients and/or high spatial frequency.

   The characteristic of clustering is particularly annoying because it greatly amplifies and localizes errors. The need for actually measuring cloud motions and observing the effect of data compression on the motion vector field for a given cloud target or targets is, perhaps, most evident here.

4. Time Lapse Display Results.

The time lapse displays of the differenced images most vividly illustrate how the data compression technique removes cloud motion information as well as adding "compression noise". At a compression factor of 4, some of the 72 x 72 pixel regions were uniformly gray, indicating that all information is preserved. Some centers of compression noise are visible. At a factor of 6, one could begin to identify small regular lineal features and edges in the Jupiter images which have higher spatial frequency and higher contrast. In low contrast areas the compression noise generated localized regions of apparent uniform flow, which might be picked up by the correlation algorithm as motions and possibly be misinterpreted. Such regions of "excess motion" are integrated in space by the human brain to show uniform motion, but the excess motions were in actuality randomly distributed in both direction and amplitude as the following residual analysis showed. The more highly compressed images showed only a larger spread in deviation from the motion of the same features measured in uncompressed images. These localized regions of motion are found most often near the centers of the compression noise.

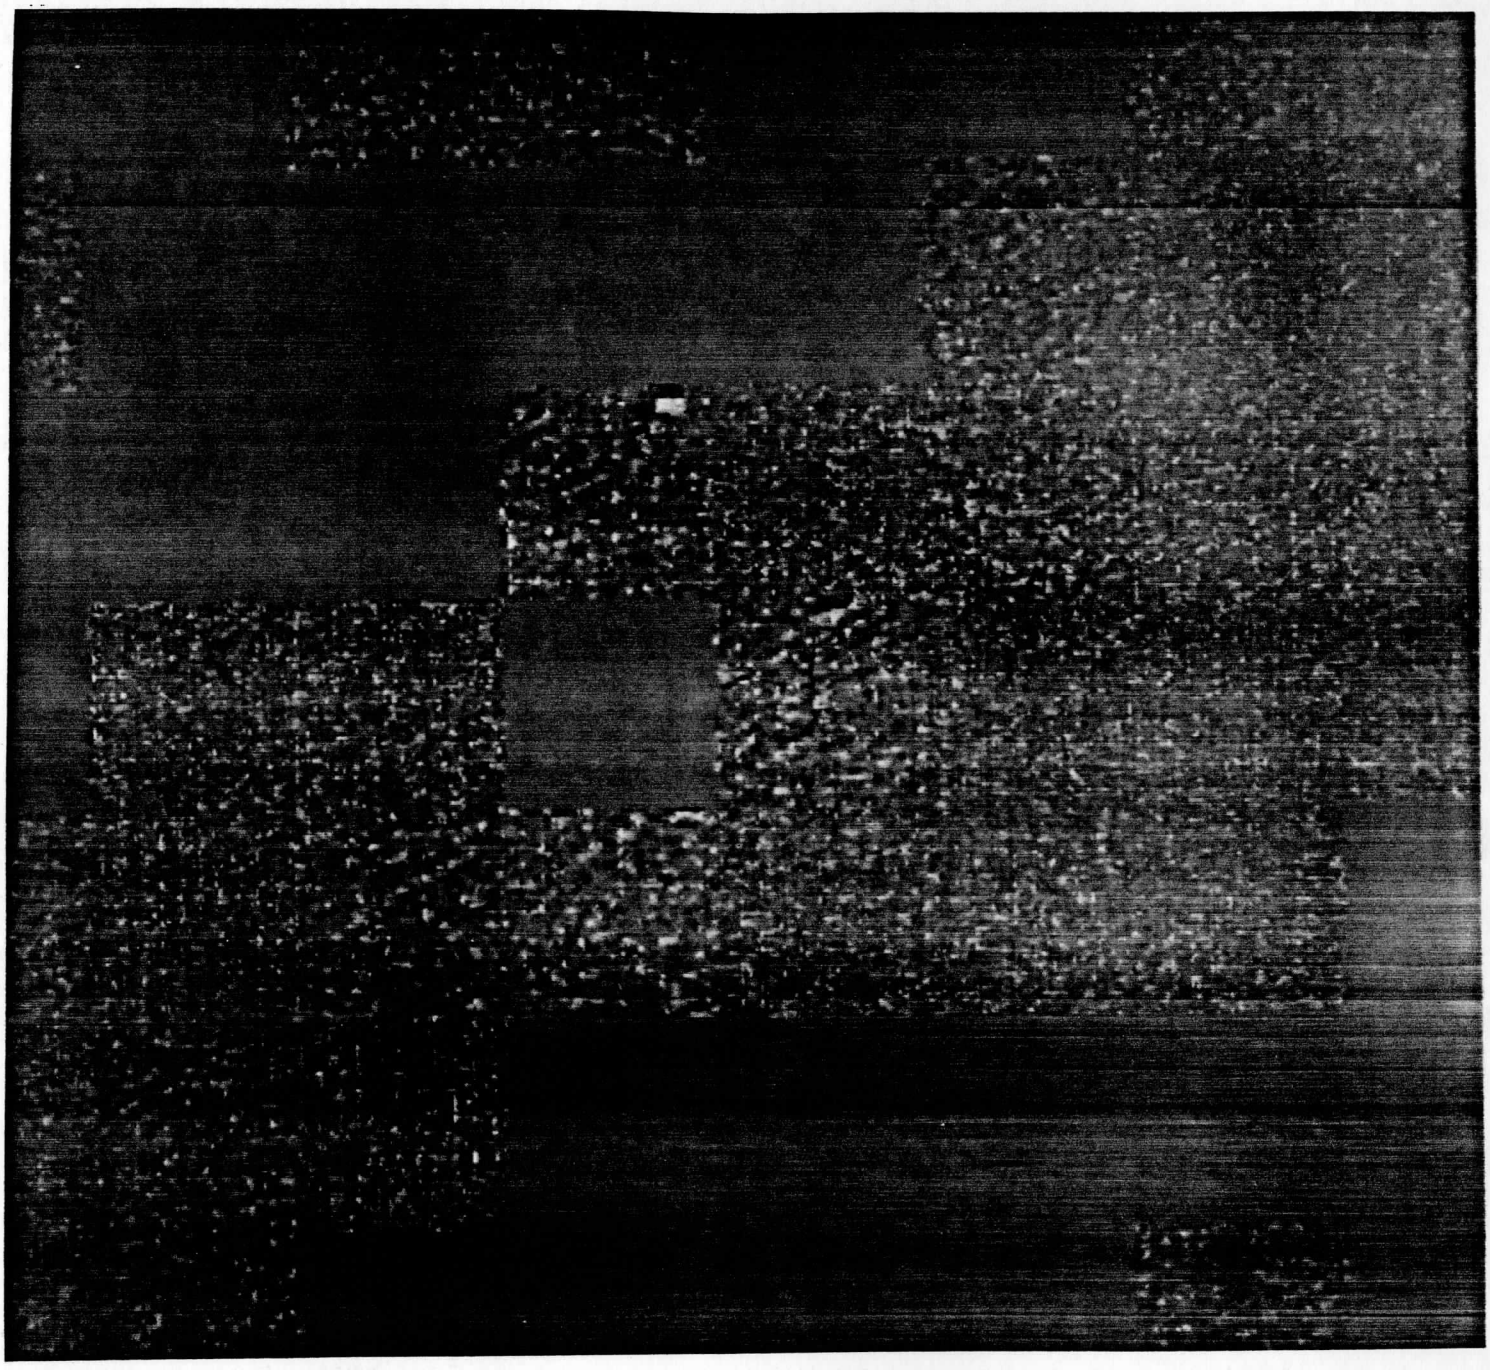

Figure 3a. Difference image showing compression noise introduced at 4:1 compression factor. Blank gray areas show regions of exact reconstruction. All the difference images have been contrast stretched to more clearly show the compression noise.

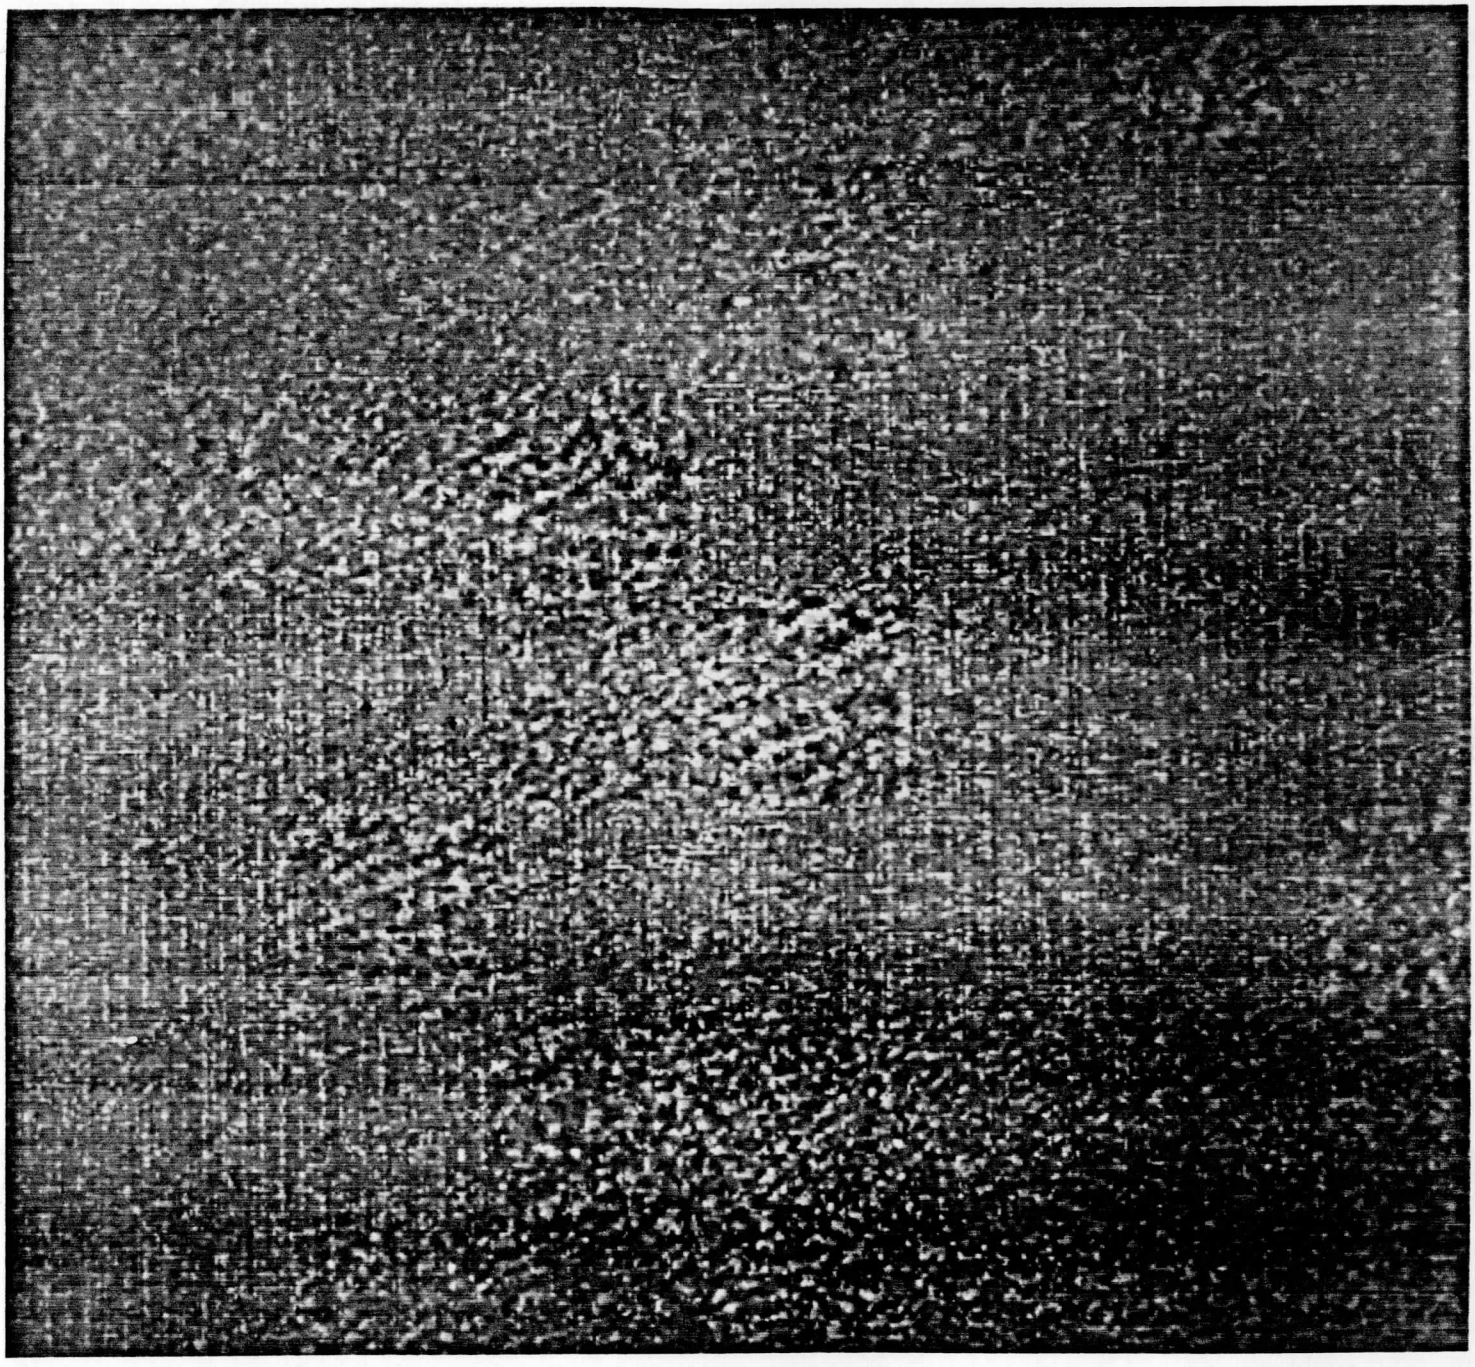

Figure 3b. Difference image showing compression noise introduced at 6:1 compression factor. The size of the grain pattern in the noisy regions shows whether 2 x 2, 4 x 4, or 8 x 8 pixel sums were edited out of the image by compression. Smaller noise grain implies better image reconstruction.

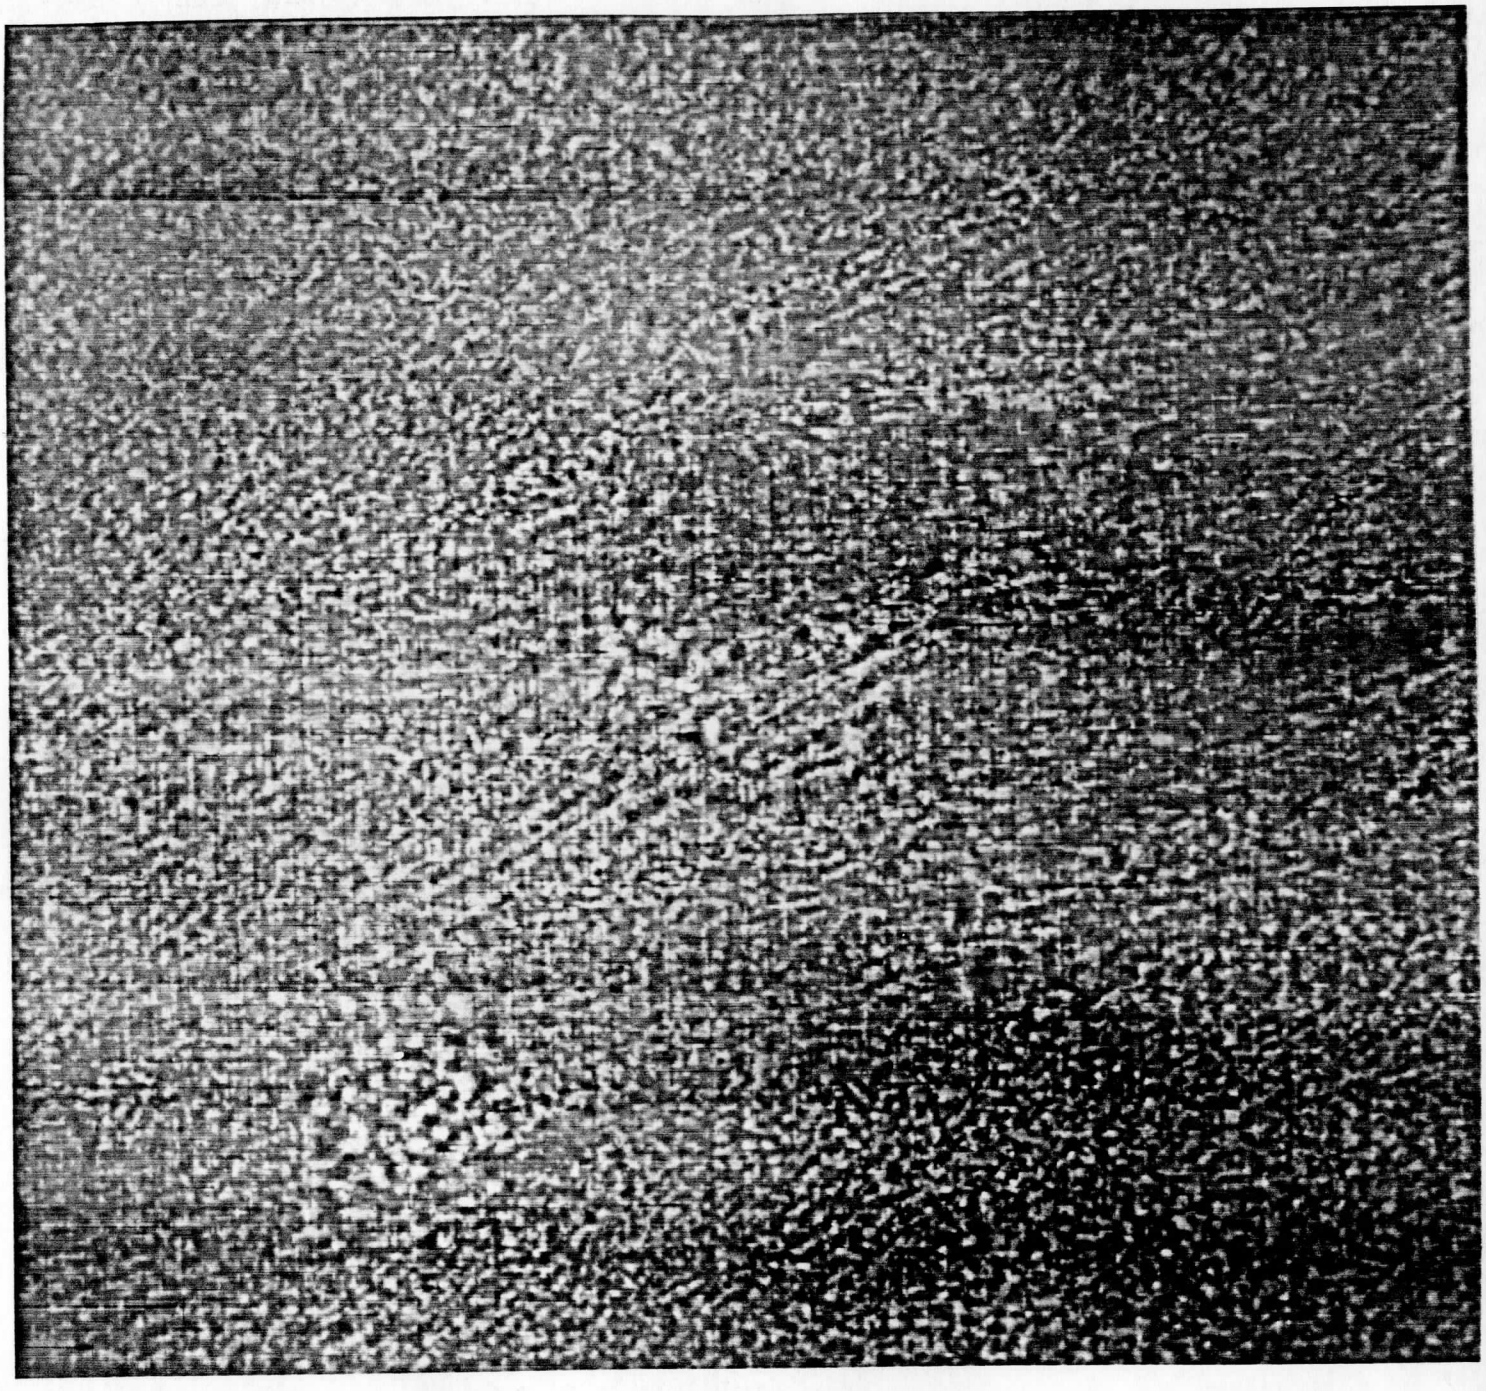

Figure 3c. Difference image showing compression noise introduced at 10:1 compression factor. Notice that some edges of cloud features are beginning to be seen in the noise pattern as the compression factor increases.

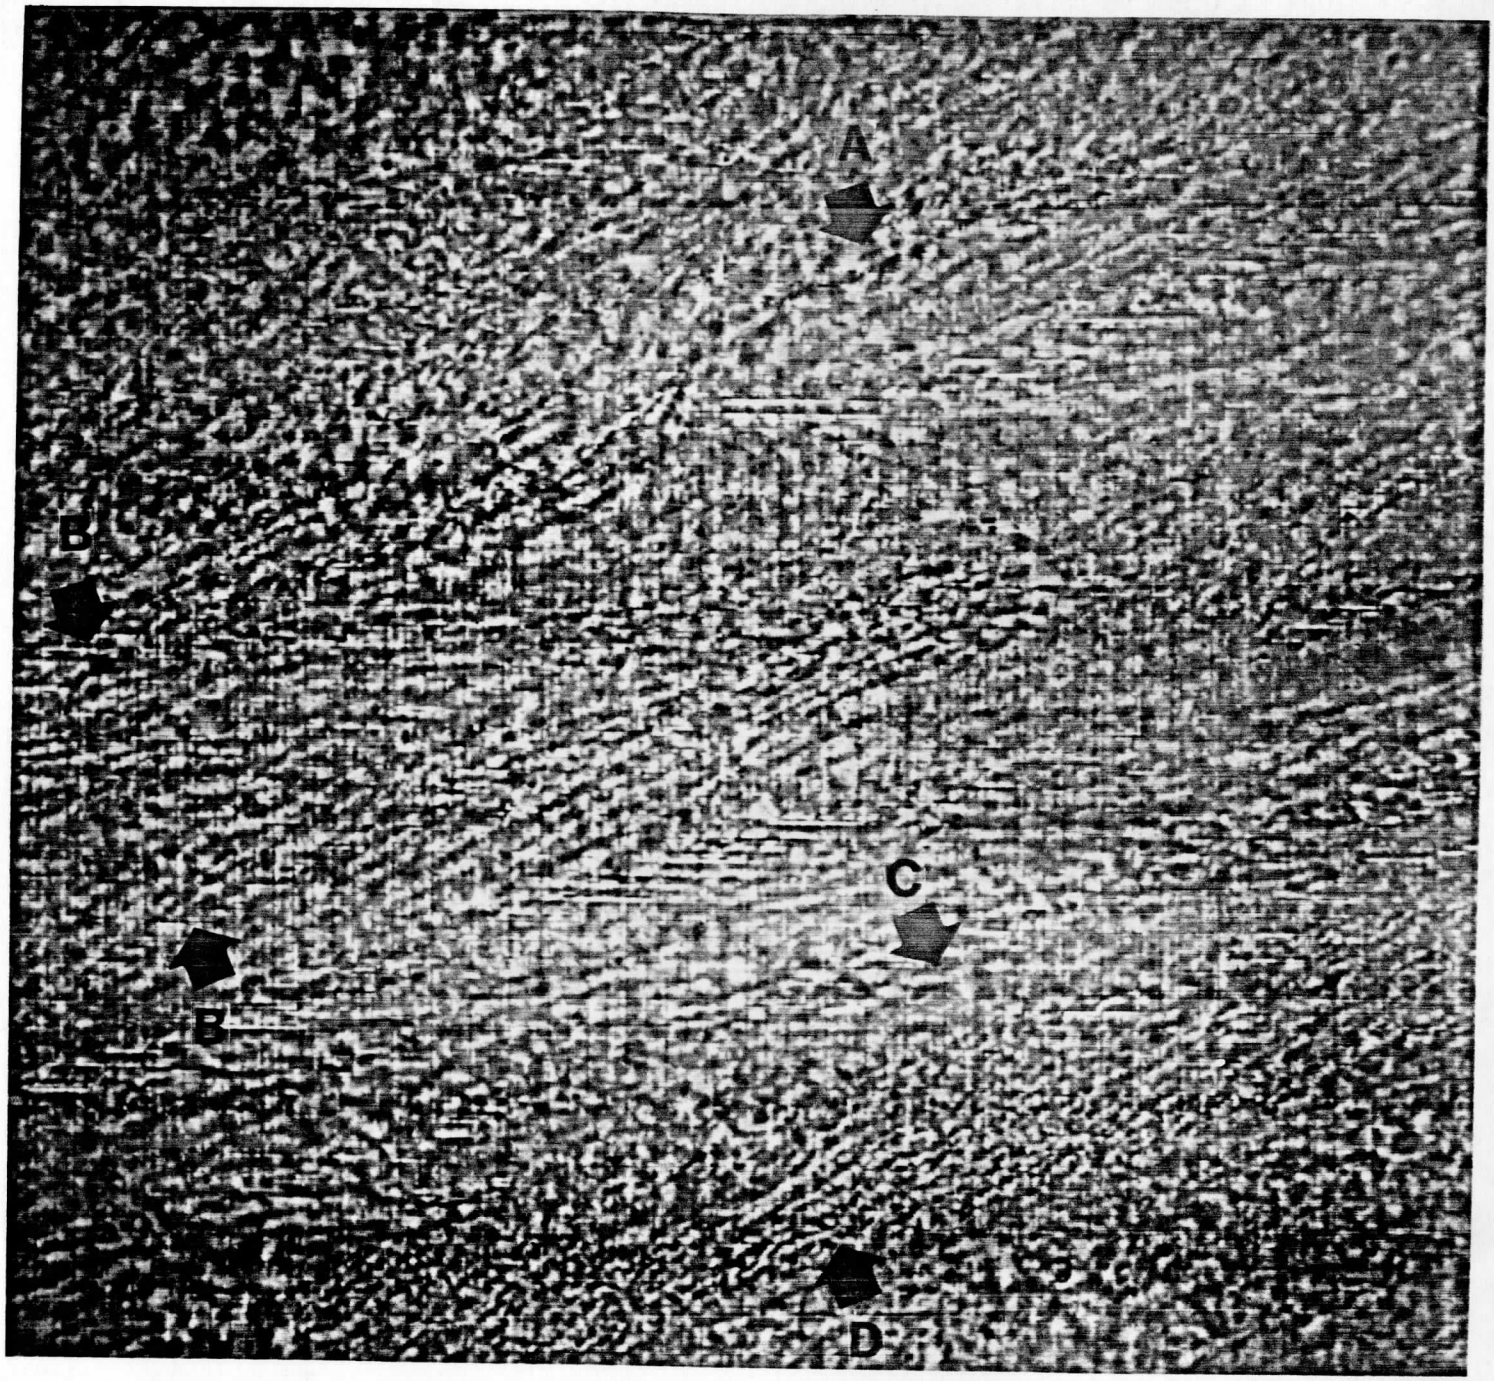

Figure 3d. Difference image showing compression noise introduced at 16:1 compression factor. Note that many lineal features are recognizable, when this image is compared with the original, shown in Figure 2a, and the 16:1 compressed image, shown in Figure 2b. This indicates that some information has been removed by the compression process.

At a compression factor of 10, the data loss was severe in both the Venus and the Jupiter images. Coherent "excess motion" fields of 3-5 DN amplitude were present throughout the images, and we also saw recognizeable true cloud motions and cloud edges representing information lost in the data compression process. At a factor of 16, the effect of compression is obviously intolerable by even crude visual standards.

We expected substantially larger errors in measurements using correlation procedures. Surprisingly, the errors were not as large as one might expect. This lack of large motion errors was suggested by the results of the quantitative histogram and cloud target brightness analysis. Apparently almost as many pixels increase in brightness as decrease, so while individual pixels may change several DN up or down, the mean value is only slightly affected in a region as large as a 15 x 15 grid, and there is sufficient information available for a successful correlation. Nevertheless, we feel obligated to point out that the amount of information removed from the cloud images at 10:1 and 16:1 leaves us reluctant to fully support such high compression factors.

V. RESULTS OF CLOUD MOTION MEASUREMENT ANALYSIS

A total of 207 cloud targets were selected for manual tracking in a triplet of Jupiter images shuttered one rotation apart (numbered FDS 15975.17, 15987.41, and 16000.05). A triplet of images is used because it permits quality control checks on both the image navigation and the cloud tracking and velocity determination procedures. Jupiter images were chosen for the wind measurement test because Jupiter has a wide variety of cloud types of varying size and contrast, and thus serves as a more stringent test of the influence of data compression on correlation methods The three green Jupiter images chosen could not be reconstructed at the BARC 2.5 compression factor, as the Venus images could. Their measured "activity" or "entropy" ranged from 2.94 bits/pixel to 3.06 bits/pixel.

The images were navigated in the usual manner, using the SEDR data ("Supplementary Experiment Data Records" containing spacecraft position and distance information) and a least squares fit was done on the bright limb points to find the planet center in the images (identified with the sub-spacecraft point). Then the 207 cloud targets were individually selected by placing a small cursor over the clouds displayed on the video/CRT terminal. The targets had to be clearly discernable in the all the images and within 60 degrees of the spacecraft subpoint. This eliminated clouds near the limb, which can be distorted due to perspective, and clouds near the terminator, which have contrast too low to be easily tracked through three successive images. Care was taken during the target selection to see that the features

had edge details in orthogonal directions so that the cross correlation tracking to follow would be more likely to succeed.

1. Navigation Quality Control

The next step was to do a quality control check on the image navigation by looking at the residual distributions of the velocities. We want to be sure the the conversions from line-element positions of clouds to latitude-longitude positions are as accurate as possible. If the planet center is accurately found in each image, no systematic errors in the motions will be apparent between time interval T1-T2 and time interval T2-T3. Misalignment of planet centers, or a wrong distance or scale factor introduced in the navigation will cause a systematic acceleration to appear in the data. This will make the mean velocities in the u and v directions different for the two time intervals, in spite of the fact that the same cloud targets are being tracked over each time interval. Moreover, such a systematic displacement or acceleration in any arbitrary direction is likely to couple the u and v accelerations in such a way that the velocity residuals between the two time intervals will be correlated.

For every vector pair for a given target we calculated the two velocity residuals

$$\Delta u = u_{T1-T2} - u_{T2-T3}$$

and

$$\Delta v = v_{T1-T2} - v_{T2-T3}$$

and the covariance

$$\Delta u \Delta v = (u_{T1-T2} - u_{T2-T3})(v_{T1-T2} - v_{T2-T3})$$

These distributions are plotted in Figures 4-6.

We also calculated the quantities

$$r = \langle \Delta u \Delta v \rangle / \sigma_u \sigma_v$$

$$\langle \Delta u \rangle = \frac{1}{N} \sum \Delta u_i$$

and

$$\langle \Delta v \rangle = \frac{1}{N} \sum \Delta v_i$$

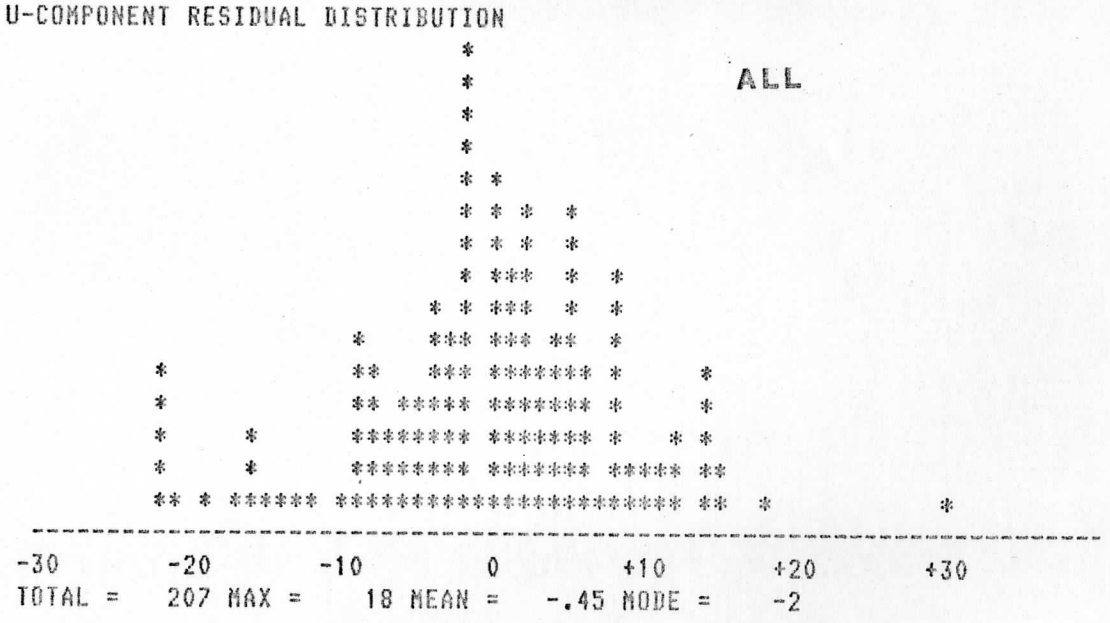

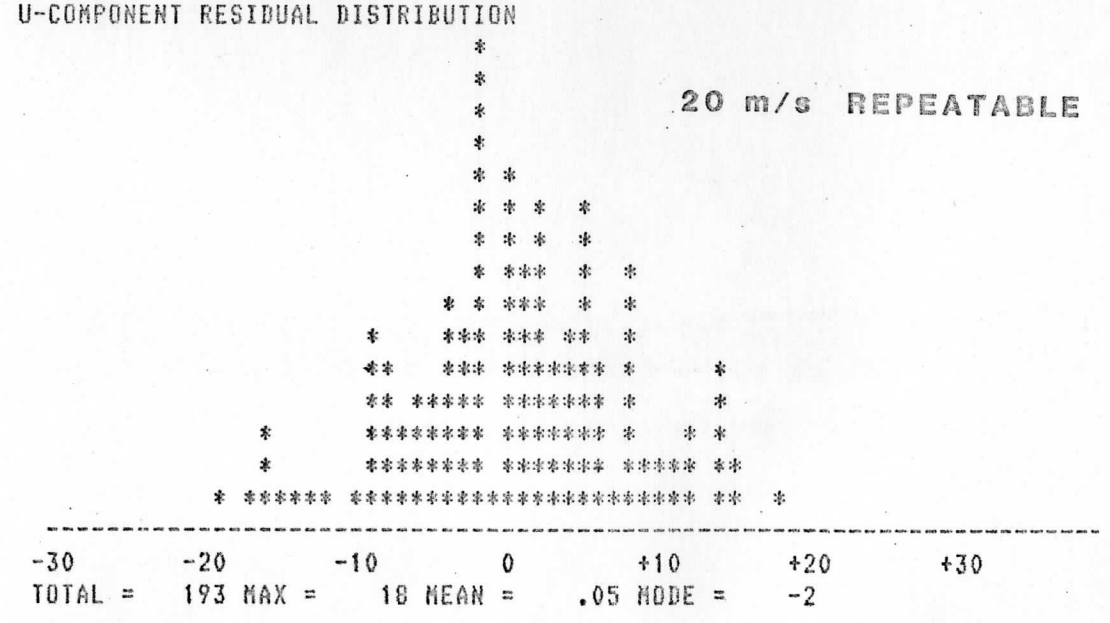

Figure 4. Residual distribution of u-component winds for the two time intervals studied.

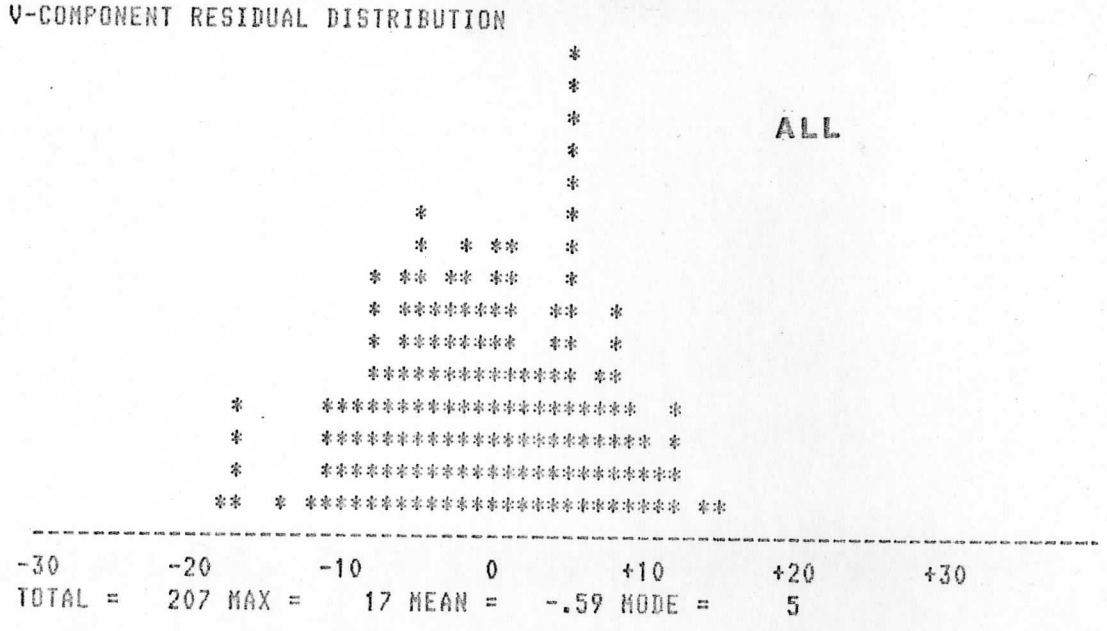

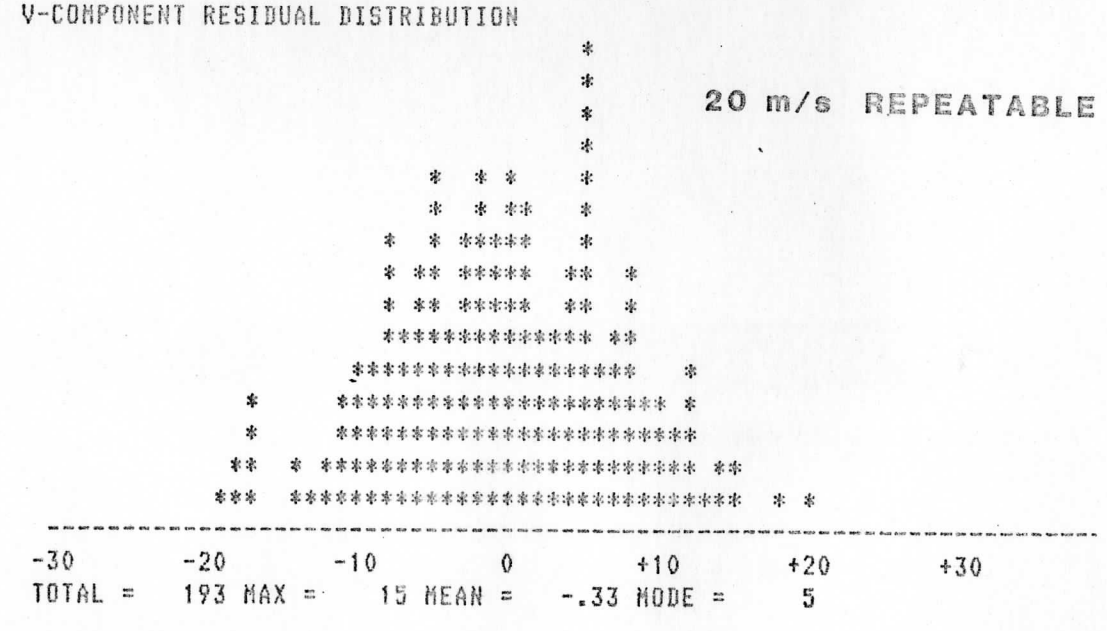

Figure 5. Residual distribution in v-component winds for the two time intervals studied.

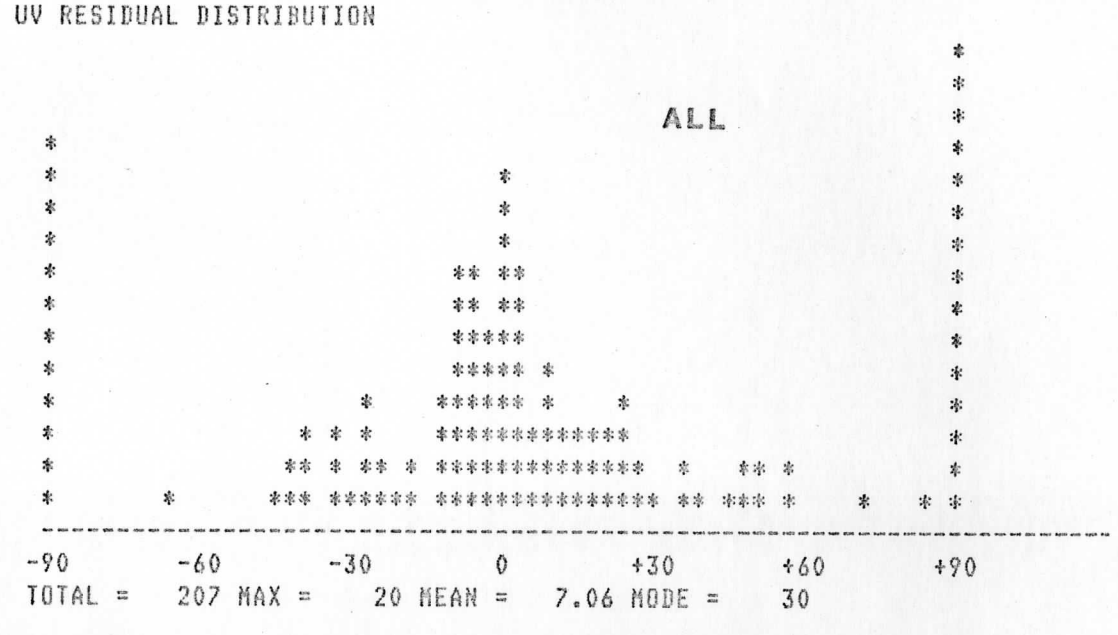

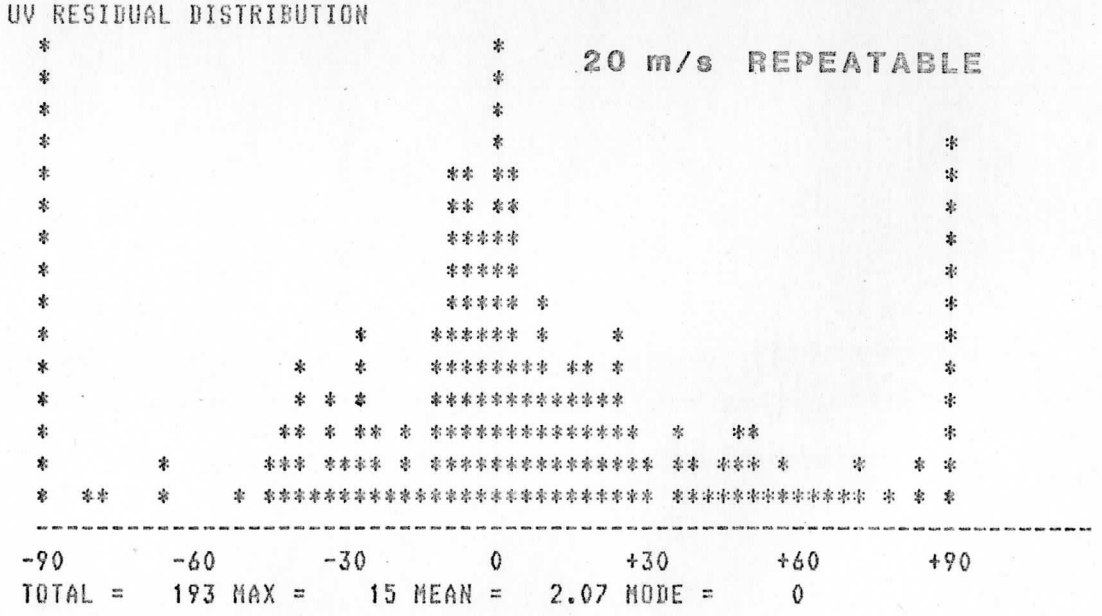

Figure 6. The product of the u and v residuals for the two time intervals studied.

The means of the residual and covariance distributions and their rms widths are displayed in Table I. For this set of images, the expected wind velocity error is ~3.5 m/s at the spacecraft subpoint on the planet for a 1 pixel tracking error or image alignment error. From the values of Δubar and Δvbar, we conclude that the images are properly aligned to <1/2 pixel, and from the width of the distributions we can estimate the rms sum of manual target tracking error and true velocity dispersion at about 2 pixels in v and almost 3 pixels in u. It is impossible to separate these two sources of dispersion using residual analysis alone, so it should be kept in mind that the dispersion represents a contribution from the storms and eddies on Jupiter as well as the effect of tracking error in the images.

Table I

NAVIGATION QUALITY CONTROL MEASURES

	ALL 207 PAIRS mean....rms	20 m/s REPEATABLE mean...rms
Δubar	-0.45 + 10.17	0.05 + 8.03
Δvbar	-0.59 + 7.63	-0.33 + 7.60
ΔuΔvbar	7.06 + 85.	2.07 + 66.
r	0.0875	0.0342

If the systematic error is zero, the velocity residuals should be uncorrelated and the resulting width of the covariance distribution ΔuΔv should be the product of the residual distribution widths Δu*Δv. This appears to be approximately true. We would not expect the correlation to be completely zero since there are some true accelerations in the vortices on Jupiter which will couple u and v. Note that all three distributions are unskewed, as they should be. Skewing would indicate the possibility of misalignment or of scale errors, making velocity residuals dependent on the velocity or position of the cloud target on the planet.

We can restrict the calculation of the correlation coefficient for Δu and Δv by adding a repeatability criterion. Since the random tracking error, determined from the distribution widths in Figures 4-6, is 6-8 m/s rms, we multiply by 3 to obtain a limit of about 20 m/s, which should be well out on the tail of the distributions. We require then, that all vectors for each

target agree to within 20 m/s over the two time intervals T1-T2 and T2-T3. This cuts the sample from 414 vectors to 386 and reduces r by 60% from 0.08 to 0.03. Examination of the vectors which failed to meet the 20 m/s repeatability criterion shows that practically every vector had a large variance in the u-component. Thus, the high zonal velocities and the presence of horizontal shear near the jets has tended to magnify the amount of correlation present in the motion data because of the fact that most of the large u-variance vectors coincidentally had v-variances of the same sign.

2. Choosing Cross Correlation Parameters

Having verified that the navigation and target selection procedures are reasonably unbiased, we next choose a target size for the cross-correlation cloud tracking. The image resolution is ~115 km/pixel. We desire the cloud targets to be <2000 km in size. Otherwise we would be likely to be tracking whole storms and not the small cloud features carried by the wind. This limits the maximum target size to about 20 pixels. A practical minimum is about 10 pixels, since if one gets too close to a resolution element in target size, it is too easy to obtain false correlations on noise. Three tests were made, using 10x10, 15x15, and 20x20 target grid sizes with a 5 pixel maximum lag. The maximum lag defines a search grid larger than the target grid, in which the correlation will be done. Since the terminal operator can define cloud target position easily to within 2 pixels, the 5 pixel range around the chosen target position is more than adequate to achieve a successful correlation if one is possible. For a lag of 5 pixels in each direction, the correlation matrix will be 10x10 pixels in size, and the search grids will be 20x20, 25x25, and 30x30 pixels, respectively.

The definition of the cross-correlation matrix element at lag position (p,q) is:

$$\rho_{12}(p,q) = \frac{\sigma_{12}(p,q)}{\sqrt{\sigma_1^2(p,q)}\ \sqrt{\sigma_2^2(p,q)}}$$

where

$$\sigma_{12}(p,q) = \sum_{i=1}^{I1} \sum_{j=1}^{I2} \{ [G1(i,j)-\overline{G1}]\ [G2(i+p-1,j+q+1)-\overline{G2}(p,q)] \}$$

and where

$$\sigma_2^2(p,q) = \sum_{i+1}^{I1} \sum_{j=1}^{I2} [G2(i+p-1,j+q-1)-\overline{G2}(p,q)]^2$$

and

$$\sigma_1^2(p,q) = \sum_{i=1}^{I1} \sum_{j=1}^{I2} [G1(i,j)-\overline{G1}]^2$$

and where

G1= target grid of dimension I1 x I2 pixels (-5<p<+5)

G2= search grid of dimension J1 x J2 pixels (-5<q<+5).

In each case, the cursor of chosen size was placed at the coordinates recorded for the manually selected target in each of the three images. Computer cross-correlation was done on the data at each cursor position for the range of three target grid sizes. Cross correlation success was defined as a peak within the correlation coefficient matrix. Peaks which were at the edge of the matrix or were well defined in only one of two orthogonal directions constituted a correlation failure. The correlation success ratios on the manually selected targets were 371/414=0.896 for the 10x10 target grid size, 370/414=0.894 for 15x15, and 369/414=0.891 for 20x20. We chose the 15x15 target grid for the rest of the data compression study because it was in the middle of the range of scale sizes we considered proper for wind tracking.

The positioning of the 15x15 cursor, and calculation of the correlation matrix, was then repeated as above for each of the 207 cloud target positions in each of the three images for each of the 5 data compression factors. Table II shows the correlation success rate for the various amounts of image data compression. Note the gradual falloff in efficiency from 2.5:1 to 16:1. Applying the 20 m/s repeatability criterion removes the worst few of the vectors, but doesn't change the correlation success rates, which appear to be independent of the repeatability criterion.

3. Analysis of Correlation Successes and Failures

Table II also classifies the successful correlations achieved for each target. The position of the cloud target is selected in the first image of a triplet and the position in the last two images of the triplet is determined by cross-correlation. Thus there are two vectors which are measured for each target, the vector for time interval T1-T2 and the vector for T2-T3. If the targets are easily discriminable by the person at the terminal, we would expect that the computer should have lit-

tle trouble correlating on them. There will be a small proportion of targets which fail in each time interval, usually because the features are changing shape or are elongated so that a well defined correlation peak is not obtained in both the line and element directions. Occasionally, the contrast deteriorates over time so that the feature is not as easily seen in the last image.

Table II

SUCCESS RATES FOR CROSS-CORRELATION CLOUD MOTION MEASUREMENT ON 207 TARGETS EACH MEASURED TWICE IN AN IMAGE TRIPLET

	COMPRESSION FACTOR					
	1:1	2.5:1	4:1	6:1	10:1	16:1
1. BOTH FAIL	3	3	5	3	1	9
2. FIRST SUCCEEDS	25	24	20	21	28	28
3. SECOND SUCCEEDS	13	17	16	15	20	19
4. BOTH SUCCEED	166	163	166	168	158	151
5. 20 m/s REPEAT	160	157	158	161	149	141
TOTAL GOOD VECTORS (2)+(3)+2X(4)	370	367	368	372	364	349

The table supports these expectations. Very few correlations fail in both time intervals. There is a small number of correlations which fail in either the first or the second time interval, with a slightly higher number of failures in the second interval T2-T3. The success rate runs between 84-89% for all targets selected, with some deterioration of the number of successes as the data compression factor increases. Applying the 20 m/s repeatability criterion removes a few more vectors from the collection of "good" vector pairs, "good" being used here in the sense of having successfully passed the correlation process. The total number of "good" vectors (not taken in repeatable pairs for each target) is shown in the last line of Table II.

We must also answer the question of whether it is possible for data compression to take a vector which normally would be bad (i.e. originally failed to give a good correlation peak in the uncompressed images) and make it look like a good vector when

motion is measured in the compressed images. To see if this possibility was significant, we divided the vectors into 5 different categories:

a) Vectors having a single random failure show up as the compression factor increased for that target.

b) Vectors having a single random success show up for that target as the compression factor increased, when the correlation in the uncompressed image failed originally.

c) Vectors having a transition from "good" to "bad" as the compression factor increased, when the measurement in the uncompressed image was successful.

d) Vectors having a transition from "bad" to "good", when the measurement in the uncompressed image was unsuccessful.

e) Vectors changing erratically from "bad" to "good" or "good" to "bad" as the compression factor increased for a given target. These vectors have more than a single transition or a single random change.

Table III shows the results of the survey. Classes a & c are not problems since they simply cause failures in correlation as the data compression factor increases. The number of successes goes down slightly. The bad classes are b, d, and e, as these contain known bad vectors, and these vectors should not be included in the data set for further analysis. If they are identified as good vectors, when measured in compressed images, they will contaminate the measurements.

We see that there are few cases where misidentification can cause a problem. Classes b & d have only 13 cases, and as long as one stays below 6:1 compression factor, the chances of falsely introducing bad data into the measurements is small. The odds increase considerably at 10:1 and 16:1, but even then are a small fraction of the original 349-370 vectors in the entire data set. There are 6 additional vectors in class e where a bad vector could be interpreted as a good one.

Clearly, the advantages of data compression as a way of getting more data back under given conditions will probably outweigh the slightly increased difficulty of having to measure more targets, or the relatively small risk of introducing bad vectors into the data set and misinterpreting them as good ones.

The remaining, and most important, question that must be answered is "How badly different are cloud motions measured in compressed images?" Table IV and the figures in Appendix A illustrate the results of our comparison study in terms of measurement accuracy. The figures in Appendix A compare the scatter in

Table III

CHANGES IN IDENTIFICATION OF GOOD OR BAD VECTORS INTRODUCED BY DATA COMPRESSION IN JUPITER IMAGES

	COMPRESSION FACTOR				
	2.5:1	4:1	6:1	10:1	16:1
A. SINGLE FAILURE	0	0	0	0	0
C. FAILURE THRESHOLD	1	0	1	9	18
B. SINGLE SUCCESS	1	0	1	3	5
D. SUCCESS THRESHOLD	0	1	1	1	5
E. GOOD AT 1:1 BUT CHANGES ERRATICALLY					16
E. BAD AT 1:1 AND CHANGES ERRATICALLY					6

residuals calculated between the uncompressed image winds and the compressed image wind vectors for each of the 5 data compression factors. Note that these residuals are different in definition from the residuals presented in Figures 4-6, and discussed in the section on navigation quality control. These residuals represent an additive dispersion in the cloud motion measurements caused by the data compression process. They are not related to tracking error or true dispersion in the wind field. The width of the distributions in Δu and Δv gradually increases as the higher data compression factors are approached.

We assume that the cross-correlation measurements in the uncompressed images are totally accurate. Then, if compression introduces a difference in velocity, that will be an additional error. A non-zero mean implies a systematic difference in velocities. There should be none. The unskewed appearance of the distributions in the figures in Appendix A testifies to the randomness of the residual distributions. We can calculate an rms value about the mean, which will represent an additive dispersion on top of the dispersion in the motion field (due to both the true winds and the measurement process) as shown in Table IV. The number of vectors compared in this way will be slightly smaller than the number of correlation successes because it is required that the vector be from a successful correlation

in both the compressed and uncompressed images to be included in the statistics.

Table IV

RMS RESIDUALS SHOWING DEVIATION OF WINDS MEASURED IN COMPRESSED IMAGES FROM WINDS MEASURED IN UNCOMPRESSED IMAGES

	COMPRESSION FACTOR					
	1:1	2.5:1	4:1	6:1	10:1	16:1
RMS u (m/s)	0.00	1.26	1.52	2.04	3.72	4.51
RMS v (m/s)	0.00	0.34	1.19	1.06	2.09	3.15
SUCCESSES	370	364	363	363	350	337

The rms differences imcrease with increasing compression factor, although the meridional velocity differences are slightly smaller for 6:1 compression than for 4:1. The trend is clearly evident, however. More important is the fact that the increase is not linear. At lower compression factors, the increase in dispersion added by compression is insignificant compared to the 6-8 m/s dispersions in the uncompressed images in Figures 4-6. At 16:1, however, they are becoming comparable. Figure 7 shows the rms sum of the data compression dispersion added to the true dispersion in the wind field and the dispersion arising from the measuring process. The curves show how the total rms measurement error from all sources climbs as a function of compression factor in the Jupiter images.

We also undertook an examination of the correlation process itself (in a statistical sense) to see if we could identify how the correlation process is affected by data compression. We studied the height of the correlation peak, comparing how it changed as the compression factor increased. Little difference was observed in the height. A drop of slightly more than 1% was observed, from 0.7921 to 0.7812. Table V shows the results for all compression factors.

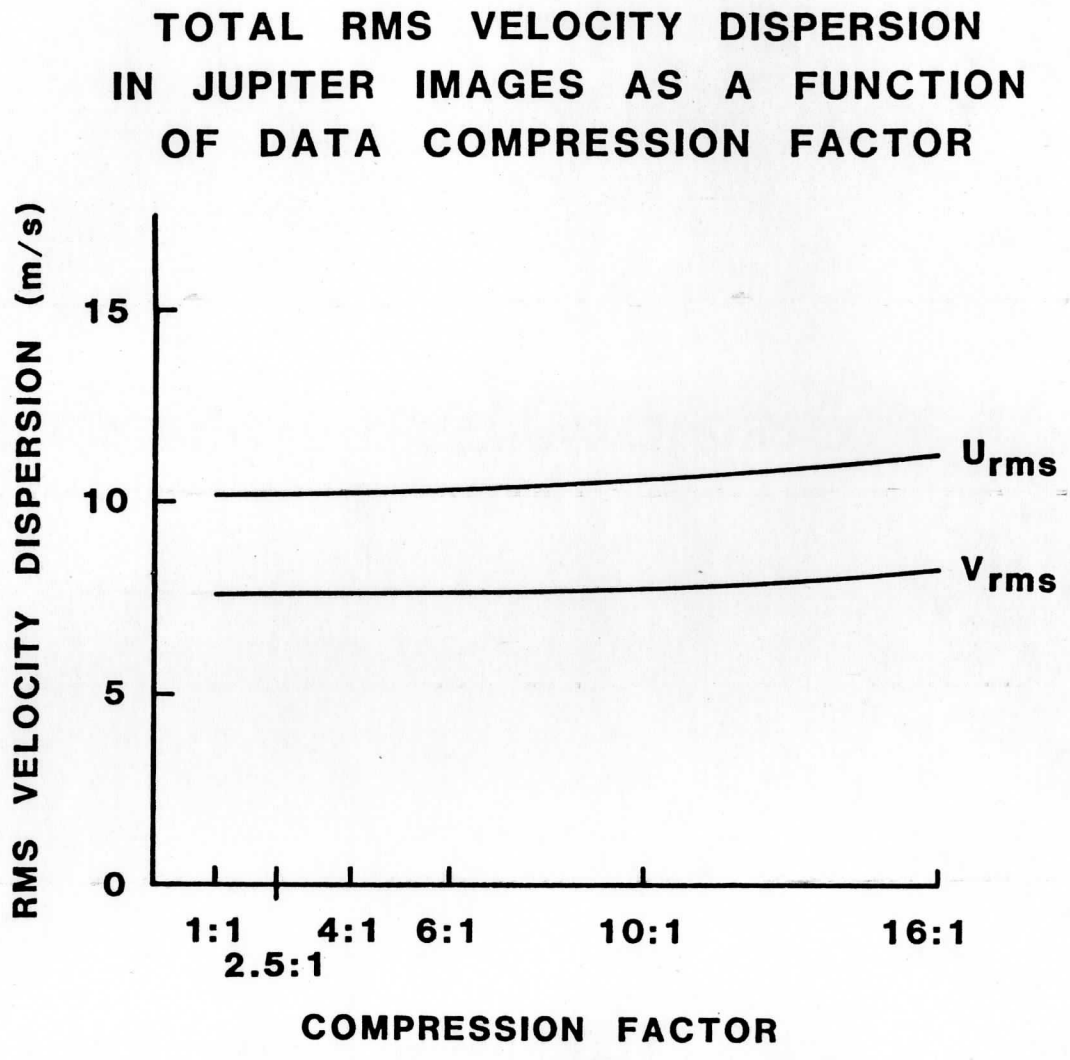

Figure 7. Shown here is the smoothed total rms velocity dispersion in the Jupiter images used in the cloud motion analysis portion of this report. Included are the effects of measurement variations and true dispersion in the wind field. The dispersion introduced by the data compression alone is shown in Table VI, and is considerably less.

Table V

MEAN VALUES OF CORRELATION PEAK AMPLITUDE, TARGET CONTRAST, AND CORRELATION QUALITY AS FUNCTIONS OF DATA COMPRESSION FACTOR (CALCULATED FOR ALL 414 VECTORS)

	COMPRESSION FACTOR					
	1:1	2.5:1	4:1	6:1	10:1	16:1
PEAK x 100	79.21	79.12	79.42	79.45	78.26	78.12
DN CONTRAST	7.04	7.04	6.96	6.91	6.81	6.48
QUALITY	-----	1.792	1.765	1.766	1.724	1.655

Target contrast was also studied, to see how the variability of the 15 x 15 cloud targets was changed by data compression. The mean target variability (expressed as rms deviation from the mean brightness) went from a high of 7.04 in the uncompressed and 2.5:1 compressed images to a low of 6.48 in the 16:1 compressed images. This change appeared to be non-linear, and closely paralleled the change in correlation success rate.

Finally we defined a cross-correlation "quality", expressed as the ratio of the correlation peak height above the mean correlation matrix value to the rms deviation about the mean. This quality measure varied for individual targets from a low of about 0.8 to a high near 3.0. Correlation quality also appears to be related to the changes in target contrast and success rate.

VI. INTERPRETATION

The scope of this effort is limited, so we cannot delve deeper into the mechanisms for the effects we see. However, it is possible to offer some qualitative explanations which could probably be developed further at some later time. The mean rms contrast of cloud features is ~7 DN of which 1 DN is the actual noise level in the Voyager cameras. If the high frequency components thrown away in the compression process represent, in sum, a total of one or two DN or less, then the signal to noise ratio for a cloud target is degraded by the ratio of DN loss to target contrast. A similar degradation occurs in the search grid. Thus, we might suspect that the correlation would be degraded by the square of the drop in signal to noise ratio. This could account for the sharper increase in sensitivity of the correlation process to compression noise as the 16:1 compression factor

is approached. Comparison of the success rates in Table II to target contrast in Table V supports this hypothesis. We might expect that the deterioration of cross-correlation success rate, depending on the image data as much as it does, would vary somewhat from image to image, as a function of not only of the planet and type of features imaged, but also of the lighting conditions which influence image contrast,and of the different spectral filters used in the cameras.

Another possibility exists. The higher compression factors result in the higher spatial frequencies being removed from the images. Consequently, the actual correlation successes likely depend on the fact that at higher compression factors, the correlation is dominated by larger scale contrasts in the 5-10 pixel spatial frequency range. Contrasts of 2-5 pixel size and amplitude ~1 DN are being mostly destroyed by correlation noise as the compression factor increases. The cloud targets must be of sufficiently large size and high contrast that lower spatial frequencies still contain a significant part of the position information.

To test this hypothesis, we went back to the measurements and organized them by latitude. We found that at 18-28 N latitude, there was a systematic increase in the correlation failure rate from 12% to 26% as the compression factor increased from 2.5 to 16. At other latitudes, no significant increase was observed. The region of Jupiter indicated by the increase in correlation failures is characterized by high spatial frequency details with a characteristic scale size of 4-6 pixels. This region is noted in Figures 2 and 3 as region "D", and the degradation of details evidenced in the difference images in Figure 3 should be noted as the compression factor increases. Verifying this hypothesis would require more statistics and some attention to the characteristics of the individual cloud targets, rather than treating them entirely on a statistical basis as we have done here.

VIII. RECOMMENDATIONS

The effects of data compression at factors of 6 or less are relatively minor and constitute no detectably serious hazard to recognition of cloud features, target tracking, and cloud motion measurement. Photometric studies in regions 15 x 15 pixels or larger would not be seriously affected either.

We cannot fully support the 10:1 and 16:1 compression factors. First, the image differencing analysis showed that a substantial portion of useful information was being removed at those compression factors. Presumably this information loss represents the 8% drop in local target contrast we observe. In unfamiliar scenes being explored for the first time, there is probably some

merit in having imaging done at lower than 10:1 compression, so the chance of losing any new information is minimized.

The tendency for degradation to accelerate at higher compression factors in those parameters we used to measure the quality of the data extraction is another concern. We cannot be sure that unknown targets and scenes will degrade at the same rate. The difficulty of tracking cloud features in low contrast images is enormous in terms of extra man-hours, eyestrain, correlation success rate, and measurement accuracy. We try very hard to stay away from the terminator for those reasons. Use of compression factors above 10:1 should probably be limited to familiar scenes whose properties are well enough known to reliably predict the types of information loss which will be suffered. The possibility of a tape recorder strategy, having compressed real-time transmission but with some uncompressed images saved for later playback, should be strongly considered. One might choose to dump to tape only those portions of image data which cannot be adequately compressed in data preserving fashion in the time available before the next frame must be exposed. It is still very attractive, however, to be able to accomodate more images in a limited amount of observation time, increasing area coverage, temporal resolution, or spectral coverage at very little added risk.

An additional possibility might be to use a more sophisticated compression algorithm which compresses different portions of an image at substantially different rates, depending on local entropy. New solid state sensors with improved storage capability and low noise need no longer be tied down to fixed frame times. The terminator or night side of a planet can be compressed at a high factor since those regions contain low contrast data. The bright side and bright limb could be compressed at relatively low factors -- almost data preserving -- to maximize information return in the parts of the image needed for navigation and containing the majority of good, useful cloud targets. Dark noise and black sky would get the highest compression factor since they have the lowest spatial frequency and are used for photometric calibration only.

It should be noted that we only navigated the original images. It was assumed that the data compression process affected the cross correlation only and not the navigation. Some of the bright limb pixels would be modified enough to cause a displacement of the limb, but there ought to have been as many displacements to the right as to the left along a scan line, so the determination of the planet center would not change substantially provided a reasonably large number of limb points was used. We used 28-29 points in each image. Nevertheless, as was pointed out in Section V.1, this is one place where a systematic error could be introduced in the wind measurement process. The fact that the residual measurement analysis shown in Table I did not show any image alignment bias to within half a pixel tends to

support the assumption that data compression would have little effect on image navigation.

In general, the concept of image data compression as a means of substantially increasing the information content of bits returned at a fixed transmission bandwidth appears to be an attractive and relatively risk-free alternative. Used with caution, and with some care given to sequencing strategies, data compression has earned our support for atmospheric studies and cloud motion measurement.

APPENDIX I

Residual Distributions of Velocity and Cross-correlation Coefficient Differences for Compressed vs. Uncompressed Images At All Compression Factors

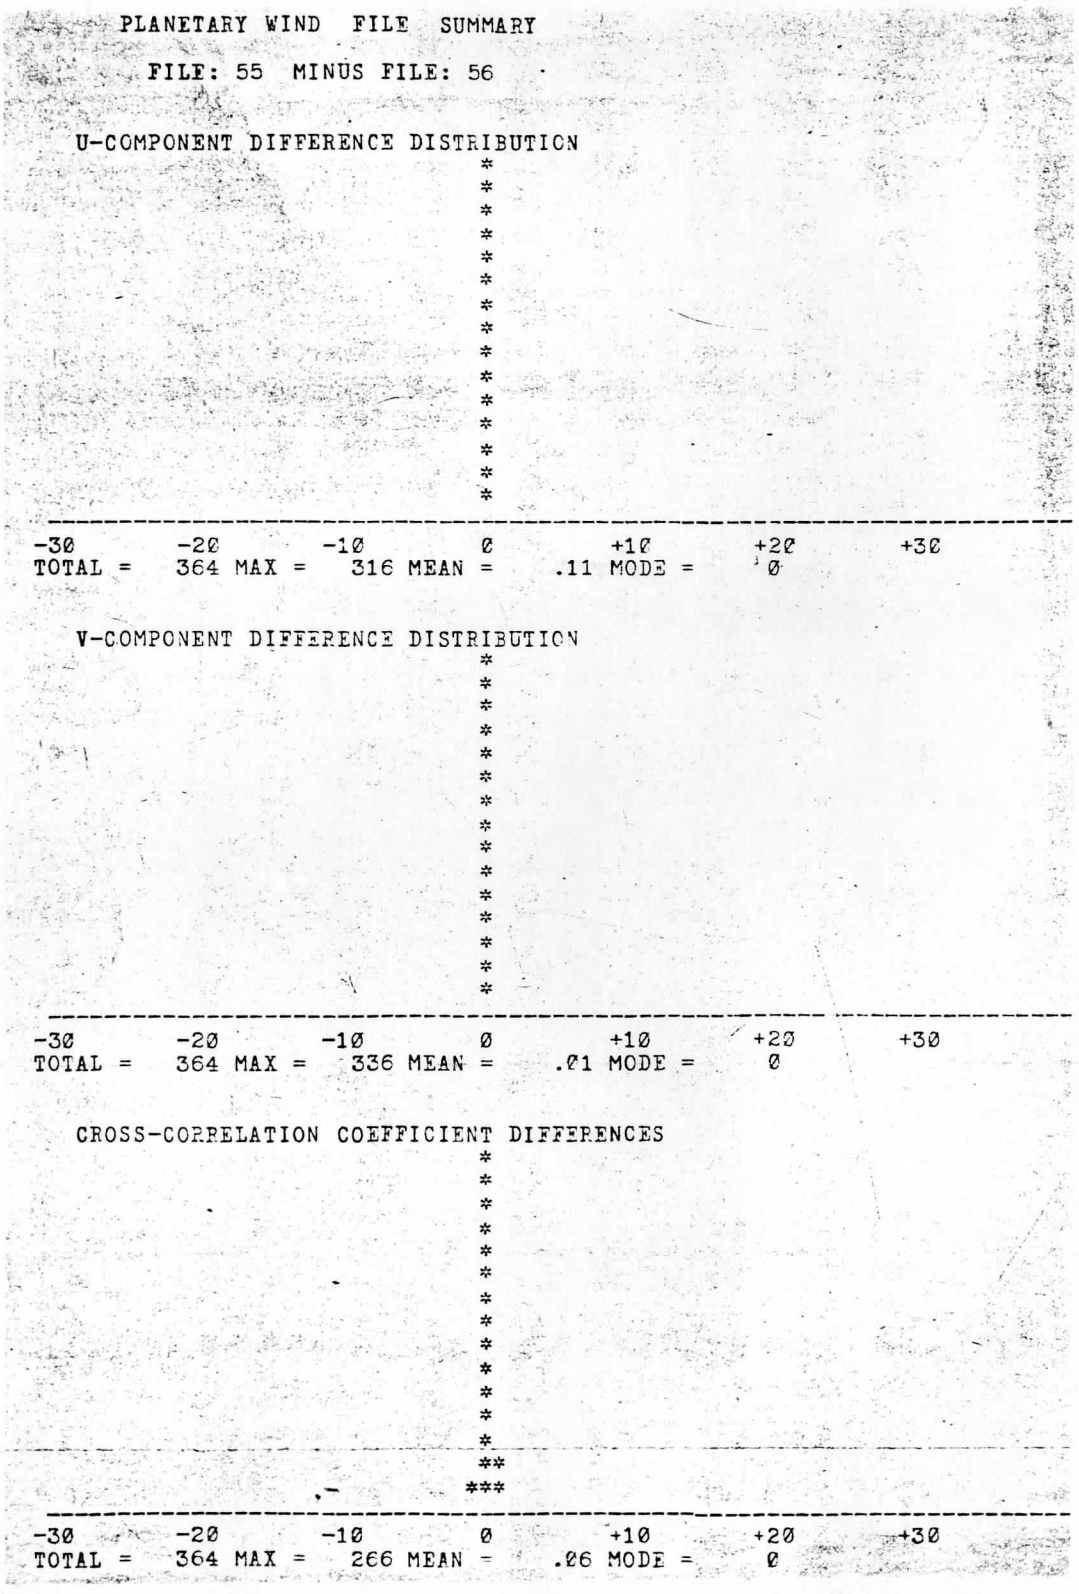

Figure A-1. Residual Distributions for 2.5:1 Compression Factor.

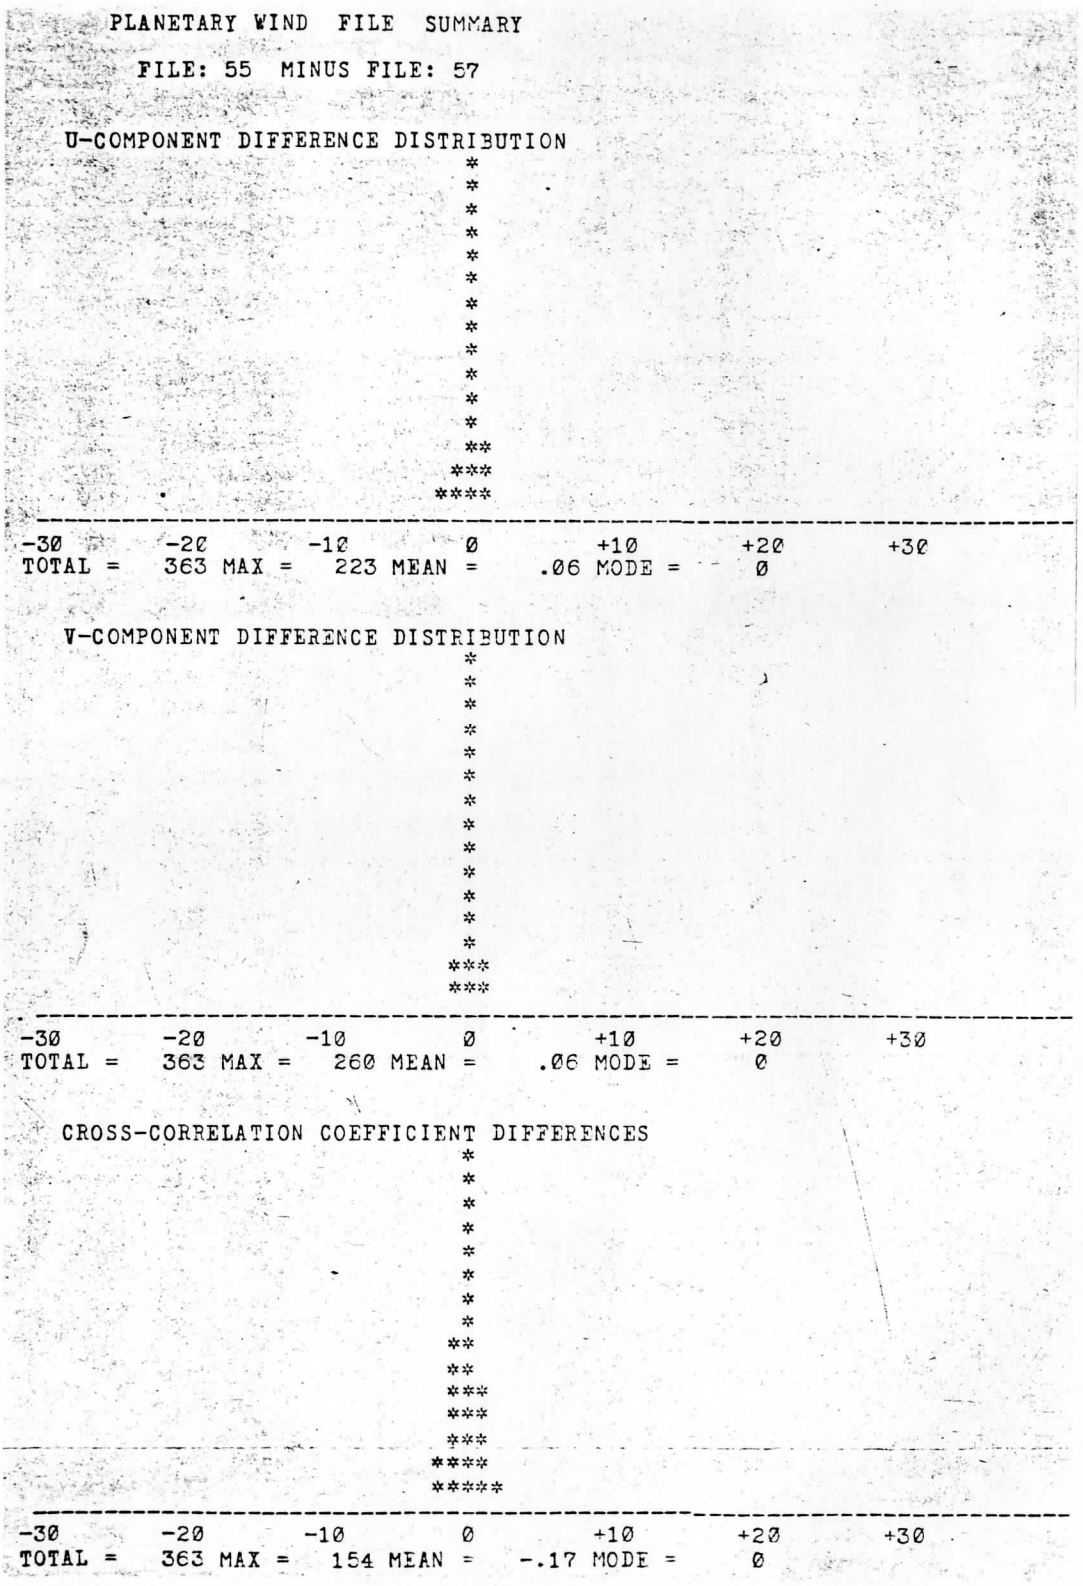

Figure A-2. Residual Distributions for 4:1 Compression Factor.

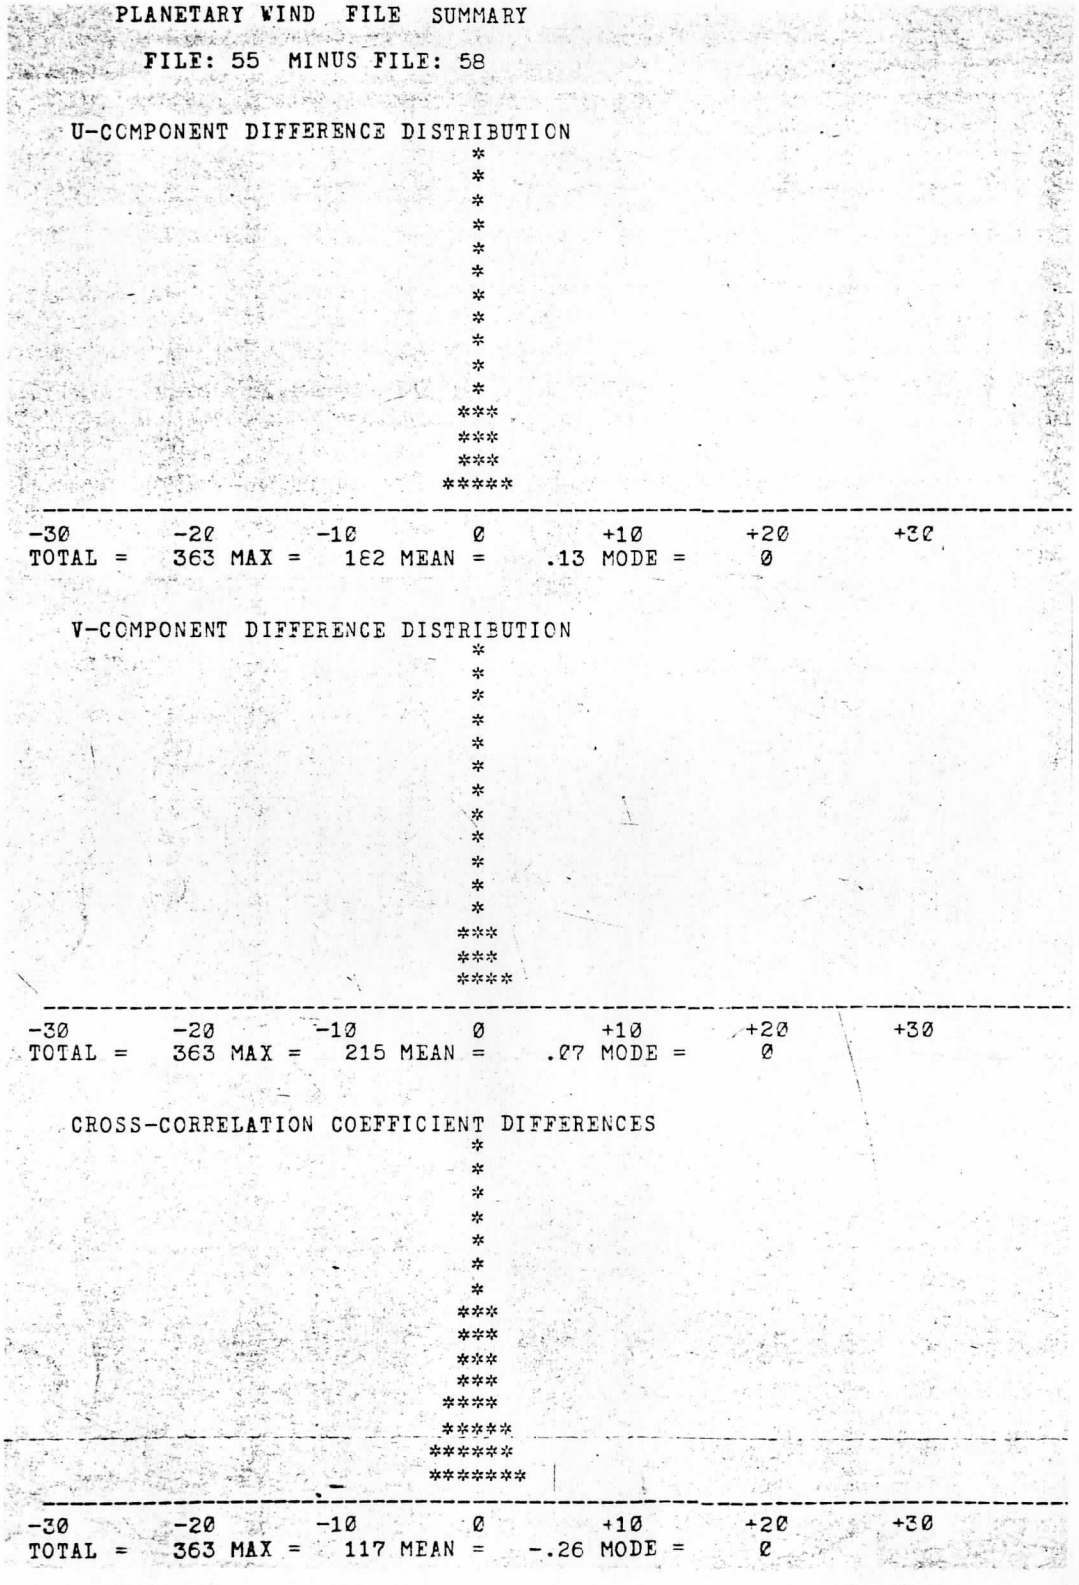

Figure A-3. Residual Distributions for 6:1 Compression Factor.

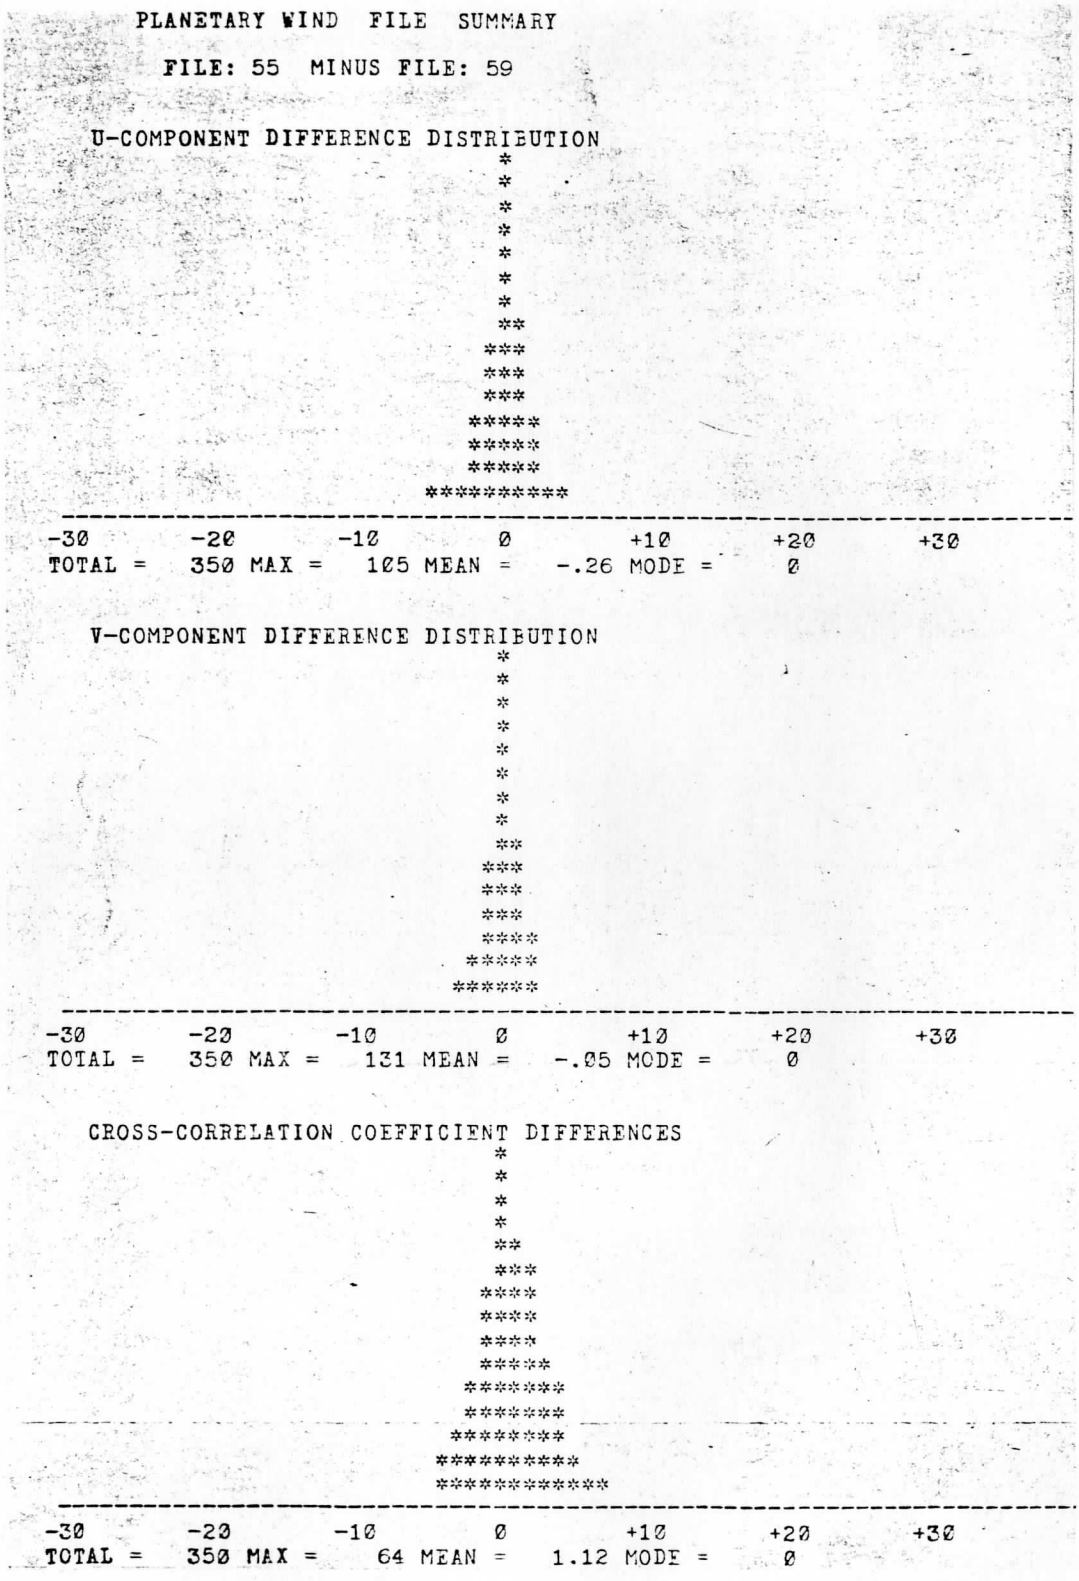

Figure A-4. Residual Distributions for 10:1 Compression Factor.

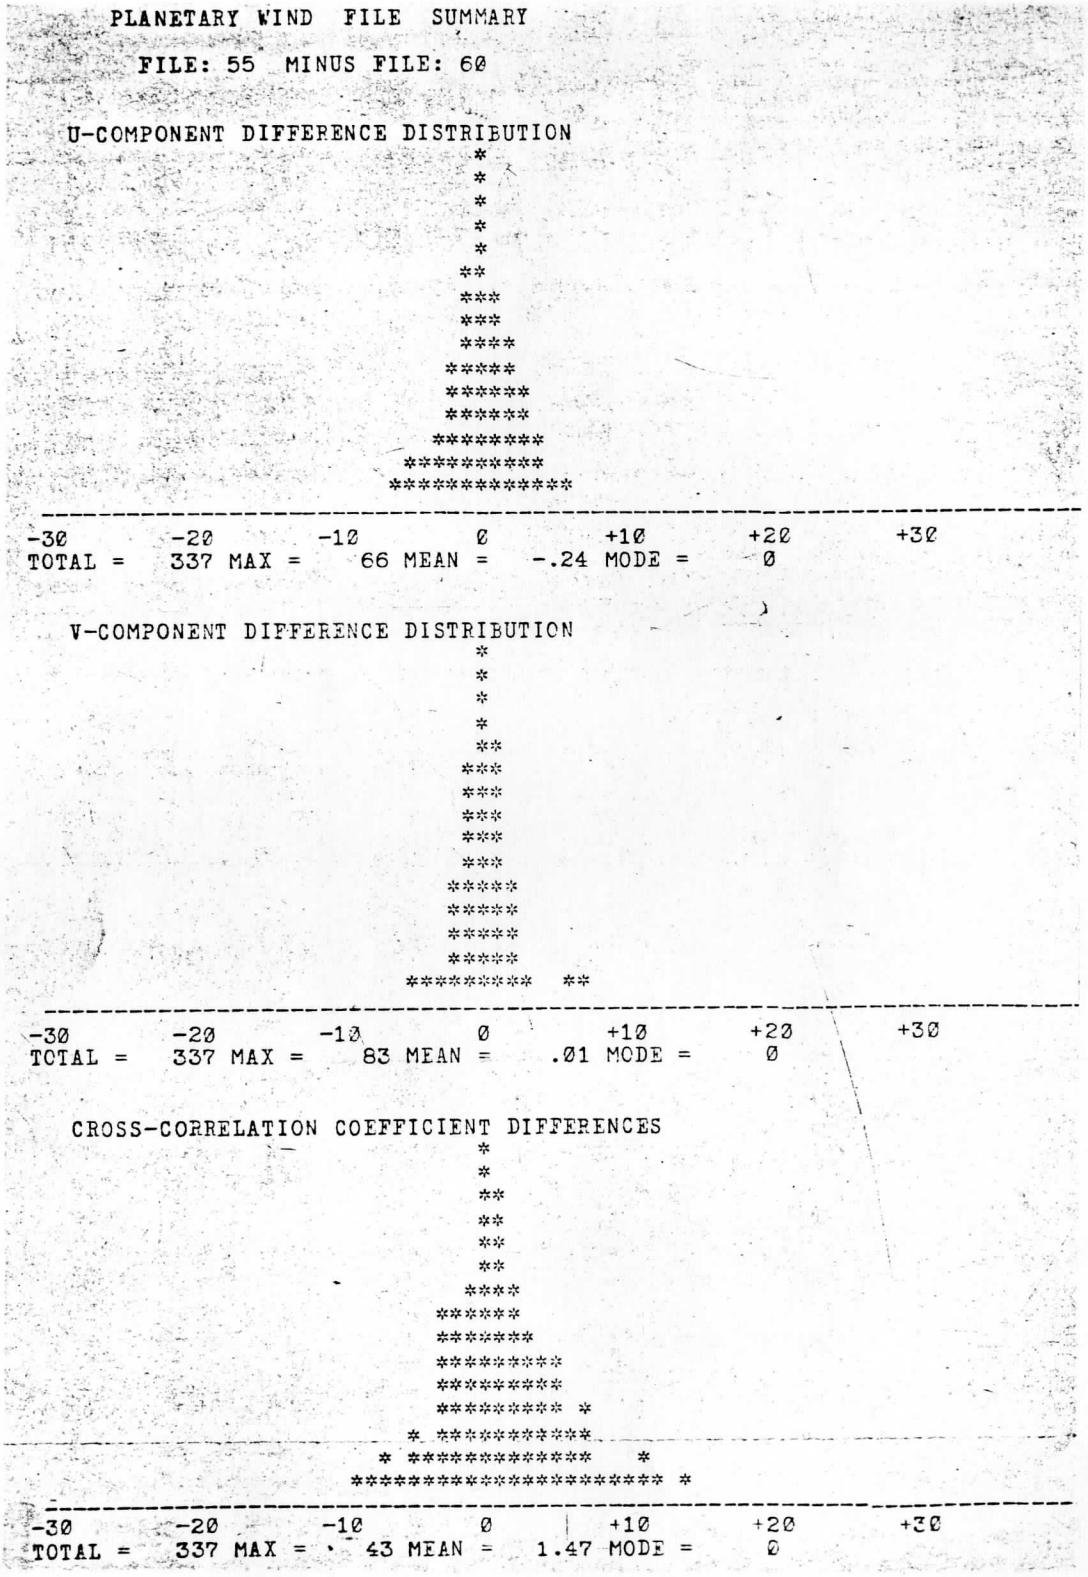

Figure A-5. Residual Distributions for 16:1 Compression Factor.