

IMAGE DATA COMPRESSION  
APPLICATION TO IMAGING SPECTROMETERS

A Progress Report

Robert F. Rice

Jun-Ji Lee

August 1981

## IMAGE DATA COMPRESSION

### INTRODUCTION

JPL has long been involved in the development of imaging data compression concepts and techniques primarily aimed at space program applications. A subset of these techniques will be key elements of the Galileo and Voyager (Uranus) imaging systems as well as in a ground based application to the transmission of IR weather data for the National Oceanic and Atmospheric Administration (NOAA). This year's efforts have focussed on investigations of the potential of these techniques to satisfy the anticipated mission requirements of Imaging Spectrometer missions as currently defined.

### Background

Noiseless Coding. Noiseless coding of data sources means compression techniques which allow exact bit for bit reconstruction of the original data. This is illustrated in Fig. 1.




Fig. 1. Noiseless Coding

Coder  $\psi[\cdot]$  maps data sequence  $\tilde{X}$  into coded sequence  $\psi[\tilde{X}]$  from which  $\tilde{X}$  can be retrieved precisely using a decoder or inverse,  $\psi^{-1}[\cdot]$ . Typically,  $\psi[\tilde{X}]$  will require many fewer bits than input sequence  $\tilde{X}$ . The average number of bits required by  $\psi[\tilde{X}]$  will vary depending on a "data activity" term called entropy.  $\psi[\tilde{X}]$  is generally considered to perform efficiently if the average number of bits required by  $\psi[\tilde{X}]$  is close to the entropy. JPL developed "universal noiseless coders" [1]-[3] which adapt to changing data statistics to ensure that such efficient performance occurs at all entropies. The general result is illustrated in Fig. 2.

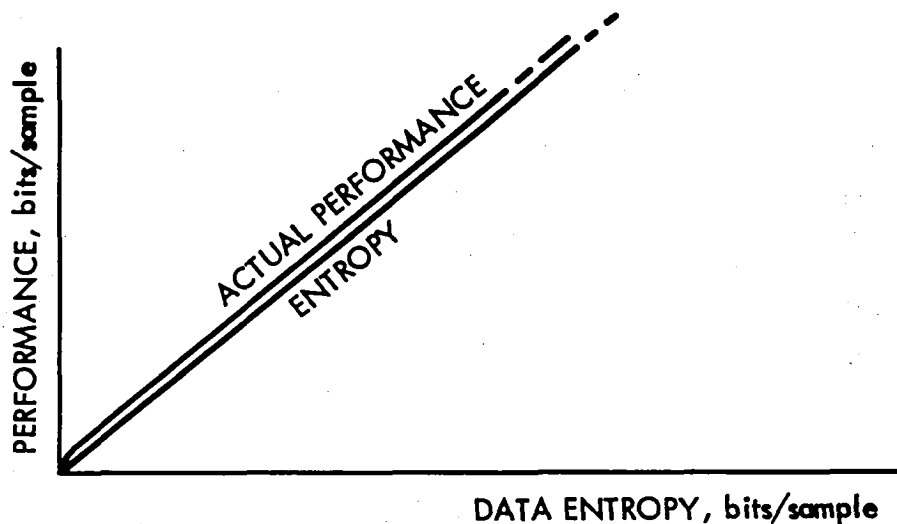


Fig. 2. Universal Noiseless Coder Performance

In image applications to deep space photos and NOAA IR weather satellite images<sup>[4]</sup> noiseless coding could be expected to reduce 8 bits/picture element (bits/pixel or b/p) by, typically, 2:1 to 4:1. However, some military reconnaissance photos required over 5 bits/pixel.

Rate Controlled Compression. The variability of the compression factor derived by noiseless coding poses an operational problem for some applications. Further, the noiseless constraint was unnecessarily restrictive in many situations, limiting the compression factors obtainable to the range noted above. These factors led to the development of rate controlled compression with the general characteristics noted in Fig. 3. A graph of image quality vs. rate in bits/pixel is shown in the figure. The two points discussed thus far are shown for no compression and noiseless coding where the abscissa (rate) for the latter is data dependent. The new feature is to be able to specify the rate: per image for a two-dimensional algorithm called RM2<sup>[5]-[6]</sup> or per line for a one-dimensional algorithm called BARC.<sup>[7]</sup> If a selected rate is above what is required for




Fig. 3. Basic Rate/Quality Tradeoff

noiseless coding ( $\approx$  entropy) then BARC or RM2 will return coded data which will allow exact reconstruction.

A gradual reduction in selected rate below what is required for noiseless coding will yield reconstructions with a corresponding gradual decrease in quality. At selected rates above 3 bits/pixel the relative performance between RM2 and BARC have generally been observed to be small, whereas at much lower rates (below 2 bits/pixel) RM2 clearly performs much better. On an absolute scale, either RM2 or BARC yield what might generally be considered "archive" quality at selected rates above 3 b/p. Here "archive" quality generally means that all objectives for which the data is intended, subjective and quantitative, can still be accomplished. If there is a constraint for archive quality or no need to achieve compression factors much beyond the 3 bits/pixel range then the simpler BARC should provide a more practical solution. The real potential advantage for RM2 and imaging spectrometer missions is the broad user controlled rate/quality tradeoff that could be applied both spatially and across multi-spectral bands to

properly emphasize features of primary interest. This is discussed further in a later section.

CCA.<sup>[8]</sup> Another development in data compression and processing at JPL which should, in some form, be applicable to the IS program is a process called the Cluster Compression Algorithm (CCA). CCA obtains data compression by working directly on multi-spectral vectors as a multi-dimensional adaptive quantizer. It has the unique property that standard supervised spectral classification procedures can be performed directly on the data in compressed form. While one might operate CCA with a 4:1 to 6:1 compression factor, the computation required to obtain classifications can be reduced by factors of 1000:1 or more. This enables such classification in real-time using low-cost computer terminals rather than special purpose hardware. While CCA has produced excellent rate/quality performance in earlier studies its most significant advantage may be on the ground as a preprocessor for classification.

#### FY81 Direction

The primary emphasis in data compression investigations for Imaging Spectrometer missions has been devoted to further development of the BARC algorithm concept. The primary motivations have already been alluded to in prior discussions and are noted below:

- 1) Link requirements for near term demonstration missions suggest that compression factors of 2 to 3:1 would adequately match instrument data rates (for primary modes) with the expected downlink capabilities.

- 2) Sophisticated adaptive modes involving user directed rate/quality tradeoffs had not been adequately studied to incorporate in demonstration missions (with high probability). Archive quality on all data would preclude user objections. The groundwork for understanding the potential usefulness of RM2 or CCA should be established, however.

- 3) Implementations at expected 300 megabit/sec input rates was a far more likely practicality for the simpler BARC.

## DATA COMPRESSION INVESTIGATIONS

This section first treats our primary efforts to improve BARC and then investigates the potential impact of a more powerful RM2 to IS missions. Time did not also permit an investigation of CCA which seems best suited for ground application.

### Efforts to Improve BARC

The original BARC algorithms are described in Ref. 7 as well as Galileo design documents. Efforts to make BARC a realizable element of a 300 Mbit/sec Imaging Spectrometer (IS) mission dealt with several issues: 1) Attempts to supplement basic BARC with sub-sample modes to extend its efficient rate-quality performance to lower per pixel rates (recall the drop off at rates much below 3 bits/pixel in Fig. 3; 2) modifications to BARC which provide more modularity and ease of parallelism needed for high rate implementations; 3) modifications which provide more accurate rate allocation procedures to enable the modularity improvement; 4) investigations of correlation detection/correction procedures to minimize the impact of channel errors on compressed data; 5) quality evaluations on data relevant to IS mission objectives.

Improvements at Low Rate. As noted earlier, if the number of bits allowed for an image line equals or exceeds the requirements of noiseless coding then a BARC coded line can be exactly reconstructed. If the number allowed is inadequate to enable such noiseless coding then the basic BARC algorithm achieves the needed reduction in bits by selectively reducing the linear quantization in blocks of 64 pixels across the line. This is a fruitful operation until the remaining entropy in a block drops below roughly 3 bits/pixel. Further reductions begin to yield contouring effects. This fact leads to the increasing disparity in performance between RM2 and BARC at low selected rates (Fig. 1). In an effort to extend the one-dimensional BARC operation to lower rates we investigated supplementing the quantization reductions with small decrements in spatial resolution as illustrated in Fig. 4.

As shown, the subsampler preceding BARC deletes 1 of  $E = n + 1$  pixels (in a staggered line by line pattern). As a result, an overall rate of  $R'$  is achieved with BARC operating at a higher rate of  $ER'$ . Missing pixels are replaced by linear interpolation.

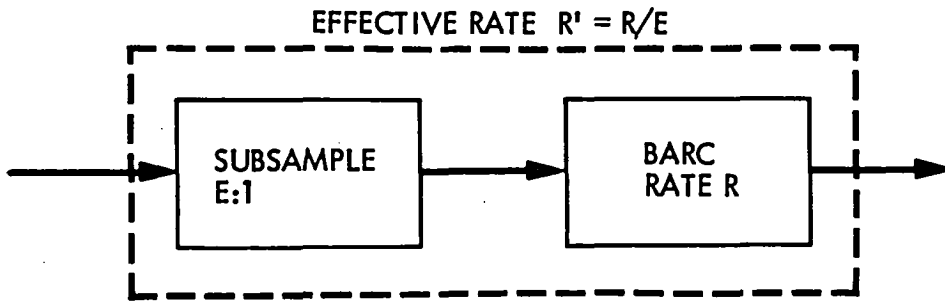


Fig. 4. BARC with Subampler

We considered  $E$  values of 1.0, 1.2, 1.25, 1.33, 1.5 and 2.0. The results are shown in Fig. 6 for two deep space Voyager images: a rather inactive Jupiter image and an active image of the satellite Callisto (we did not at the time have data which more closely modeled IS characteristics). The graph plots selected rate in bits/pixel vs. root-mean-squared-error (rmse).<sup>†</sup>

For the less active Jupiter image the subsampling modes were quite effective in lowering the rmse of BARC at per pixel rates of 2 and less. However, the combination was still in general much less effective than RM2, particularly at rates of 1 bit/pixel and less. Subsampling provided little advantage to the more active Callisto image which is closer in characteristic to IS data. The conclusions are then:

- a) Subsampling could improve BARC performance at low rates for some applications if an adaptive mechanism for selecting subsample modes can be found.
- b) It probably is of no value to IS missions.

Improved modularity, BARC2. The basic BARC performs quantization reductions over blocks of 64 pixels whereas the noiseless coding functions are performed over 16. This was partly to minimize the overhead associated with identifying the choice of quantization as well as with assuring adequate operation of the procedure for determining how quantization is determined in each block.<sup>[7]</sup>

<sup>†</sup>The standard deviation of the pixel errors in data numbers, DN.




Fig. 6. BARC/Subsample RMSE vs Rate



Experience in the Galileo design suggested that performing both quantization reductions and noiseless coding functions over the same block of 16 pixels could effect simplifications in implementation, particularly in high rate applications such as IS missions. A new method of determining block quantizations needed to be devised to avoid excessive variance in the actual number of bits used over a prescribed data span (e.g., a line). The solution has turned out to be not only simpler and actually more accurate than the old method used on blocks of 64. Compression performance is essentially the same for either approach.

Henceforth, we will assume this new approach and call the overall algorithm BARC2.

Error Protection. The effect of communication errors on compressed data such as from RM2, BARC2 or noiseless coding is much more severe than on uncompressed data. For BARC2 or pure noiseless coding the effect is a propagation of the error across a line until interrupted by a known sync word or other mechanism for a restart. For RM2 the loss may be several two-dimensional blocks of data until restart. If error events are very infrequent then the real damage is minimal. The solution for deep space is the implementation of a concatenated Reed-Solomon/convolutional-Viterbi channel. [5],[9]-[11] This channel produces "virtually error free" ( $<10^{-10}$ ) communication at the same link signal-to-noise ratio that an uncoded link would have error rates higher than 1/100. NASA is considering the application of these same Reed-Solomon codes to high rate earth orbit communications. However, the exact characteristics of communication links for IS missions is not totally defined at this time. Hence we considered the question of what could be done to accomodate errors when they occurred, however rare. Whether further error protection coding would be desirable could be answered at another time.

We first developed a normalized measure of correlation between lines (and defined for each pixel) which tends to hover about unity for the data sets we have thus far investigated. Applying this measure to decompressed BARC2 data will yield similar results, unless a communication error occurs. In the latter case, the correlation measure will eventually diverge, usually very quickly, thus detecting an error event.

The simplest use of this error detection information is the replacement of subsequent data (up to the next restart) with either the line immediately above or or perhaps the average of the line above and the line below. When error events are extremely rare this is probably a completely adequate solution. However, by projecting back to the point where the correlation began to diverge one can consider correcting the error by alternately changing bits and preceding forward until the correlation measure no longer diverges. These observations are illustrated in Fig. 7.




Fig. 7. Correlation Detection/Correction

We are still in the process of developing the error correction procedures. However, an example of automatic line replacement resulting from error detection is shown in Fig. 8. The image on the left is the result of an error rate of roughly  $P_e = 10^{-5}$  affecting an image compressed 2:1. The right hand image is the result after detection and replacement.

We anticipate that a sophisticated combination of detection/correction/replacement should enable effective BARC2 operation at much higher error rates but are focussing on the simpler task of  $P_e \leq 10^{-5}$  which probably bounds IS mission requirements.




Fig. 8. Detection/Replacement

Application to 15m Multi-spectral Data. Band 4 of a selected aircraft derived 7-band multi-spectral image of Ventura County, CA., is shown in Fig. 9. The 2410 pixel by 956 line, 8 bits/pixel image shown provides 15 meter resolution and covers the spectral range of 0.52  $\mu\text{m}$  to 12.5  $\mu\text{m}$ . Band 4 represents the range from 1.0  $\mu\text{m}$  to 1.3  $\mu\text{m}$ . The four regions enclosed by white borders are 512 x 512 subsets selected for computer processing.

BARC2 runs at 3.24 bits/pixel (or 2.5:1) on bands 1-7, subset 2, produced an overall rmse of 1.79. This is a quite accurate representation for an image with an average band entropy of just over 5 bits/pixel and a standard deviation about the data mean of nearly 30. The pictorial representation of these results is given in Fig. 10. The image in the upper left is the original Band 4 of 512 x 512 Subset 2 in Fig. 9 whereas the upper right is the BARC2 result at 3.24 bits/pixel (2.5:1). It is impossible to see the difference. The 128 x 128 bordered region in both images is severely blown up directly below. These results should support the "archive quality" assumption made earlier.




Fig. 9. Ventura County, Band 4

#### RM2 Application to 15m Ventura

The RM2 algorithm noted earlier was also run on the high resolution multi-spectral image of Ventura County, CA. The results are illustrated in the pictures of Fig. 11-13 and the rate vs. rmse graphs in Fig. 14. The images deal again with Band 4, Subset 2 whereas the rmse plots are composites over all 7 bands.

An original 8 bit/pixel (b/p), 512 x 512 rendition of Subset 2 (Fig. 9), band 4 is shown in Fig. 11 along with RM2 results at 2.0 b/p (4:1), 1.33 b/p (6:1) and 1.0 b/p (8:1). This is followed in Fig. 12 with the corresponding "diff-pics" which display the error between the original and compressed images shifted to an average value of 128. Selected blowups at 256 x 256 appear in Fig. 13.

Degradation is difficult to observe in the more realistic 512 x 512 displays for rates down to 1.33 b/p. Image quality is quite good but probably not of




Fig. 10. BARC2 on Ventura, Band 4, Subset 2




Fig. 11. RM2 on Subset 2, Band 4




Fig. 12. RM2 Diff Pics on Subset 2, Band 4




Fig. 13. RM2 256 x 256 Blowups






Fig. 14. RM2 rmse results, Subset 2

archive quality at 2.0 b/p. Note that rmse data points in Fig. 14 lie along a straight line which roughly extrapolates back to the data entropy, which for band 4 is 5.03 b/p (a completely random field would have an entropy of 8 b/p). This extrapolated graph will lie slightly below the data point for BARC2 at 3.24 b/p (and illustrated in the pictures of Fig. 10). The results are consistent with the introductory discussions.

It is important to note that the rate/quality tradeoff evidenced by these images and rmse graphs is a continuous one and is not limited to only the selected sample points presented here. A user could arbitrarily tradeoff spatial quality for spectral coverage in virtually any combination, and change his options at any time.

RM2 Global Rate Allocation. Additional gains are potentially possible if one looks further into the rate control structure of RM2. RM2 partitions a single band image into subpictures of size 32 x 32 or 64 x 64, and determines an activity measure for each. This results in an array of activity numbers as shown in Fig. 15 for an N subpicture image. The activity numbers reflect the relative need of each small area for bits. Given B total bits to use in an image an RM2 rate allocation procedure utilizes this information to designate  $B_i$  bits to be used to code the  $i^{\text{th}}$  subpicture (where  $i = 1, 2, \dots, N$ ) so that<sup>†</sup>

$$\sum_{i=1}^N B_i = B \quad \text{where } B \text{ arbitrary} \quad (1)$$

Subpictures with larger activity measures tend to receive more bits than the less active subpictures. Such allocations could in principle be applied across spectral bands or even sequences of images.

Geographic Boost. This global rate allocation procedure thus focusses bits and hence quality towards regions based on a natural measure of activity. The next potential step in performance might be obtained by simply fooling the rate allocator. Suppose for example that a certain geographic area of say 256 x

---

<sup>†</sup>A subpicture to be completely edited (not sent) receives zero bits by this procedure.




Fig. 15. Rate Allocation

256 size within a very large image (e.g., 2048 x 2048) was known to be of critical importance to a user. He desired 7 full bands at very high quality in this region but the link could support all 7 bands only at an average of 1 bit/pixel. If the position of this region could be (roughly) designated then the user's problems are easily solved. By artificially boosting the natural activity numbers,  $A_i$ , of the subpictures in the designated region, the rate allocator would unknowingly allocate a greater number of bits (and hence quality) to those subpictures at the expense of the remaining regions. Since the selected region is only  $1/64^{\text{th}}$  of the total this would have little impact on remaining regions. In general, the amount of boost would reflect the relative importance and perhaps size of the selected region over other areas. This is illustrated in Fig. 16.

Auto Boost. More generally, artificial boosts to activities might be determined by independent pattern recognition devices which looked for special features of interest. When certain features were found to be present in a subpicture  $i$  the

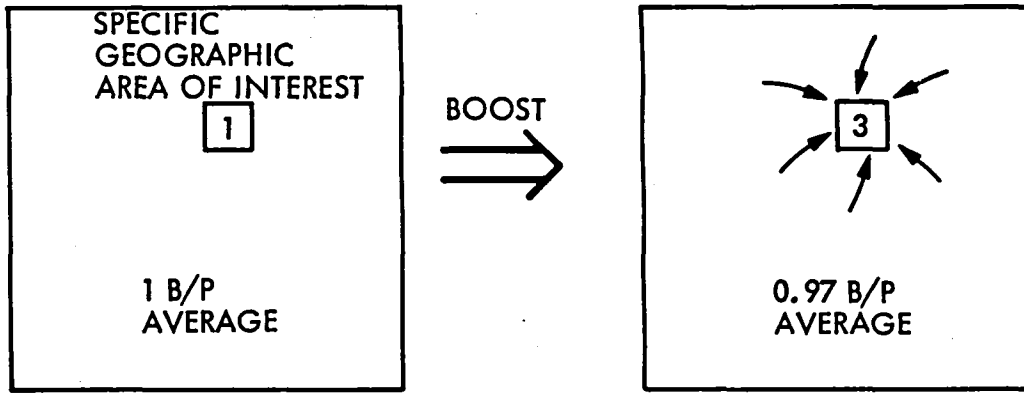


Fig. 16. Geographic Boost to Acitvity ( $\Rightarrow$  quality)

device(s) would call for a boost to the natural activity measure  $A_i$  leading to additional bits and hence quality into subpicture  $i$ .

These considerations are summarized in Fig. 17 where  $\beta_i \geq 1$  is the user generated geographic boost associated with subpicture  $i$  (a user would not pinpoint subpictures but whole regions), and  $\gamma_i \geq 1$  is the corresponding artificial boost determined by some pattern recognition device. Note that these two forms of added direction to a limited number of bits has no impact on RM2 itself and could be developed or considered as later supplements.

Principal Components. The following investigation is incomplete but offers the possibility of better performance at low rates. We now preceed RM2 with a spectral transformation which maps each multi-spectral pixel onto a new set of basis vectors which are the eigenvectors of the data covariance matrix. Such a transformation is called "principal components," or Karhunen-Loève (K&L) and other names as well. We will stick with the name principal components.

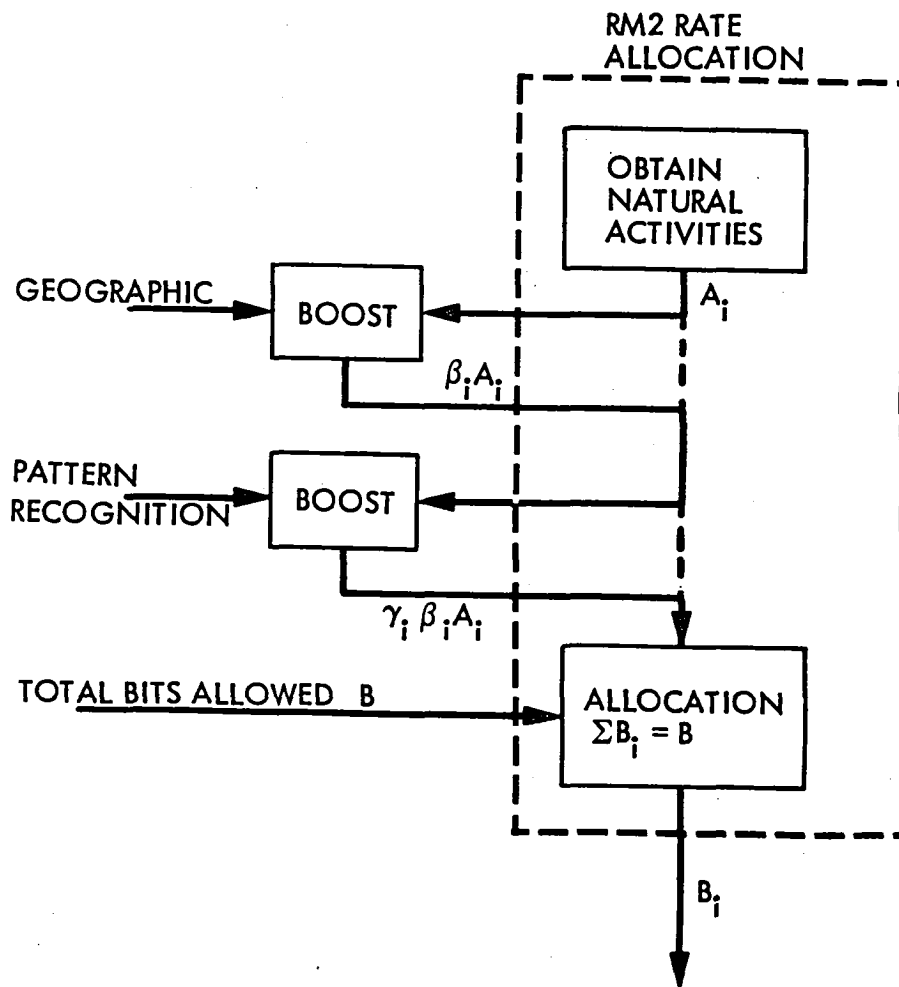


Fig. 17. Generalized Global Rate Allocation

Ideally the result of this transformation is a set of new spectral bands which are no longer correlated and for which most of the image energy has been concentrated in only a few of the new bands. The real hope for improved compression is that principal components makes use of band-to-band correlation which we have so far ignored. Ready and Wintz<sup>[12]</sup> obtained some encouraging results using less sophisticated spatial compression than RM2. However, their test sets were significantly less active than the Ventura image, Fig. 9.

RM2 Conclusions. It would seem clear that the RM2 concept offers the potential for significant increases in real information transfer rates. This would come from sophisticated user and autonomous direction of the rate/quality trade-off provided by the RM2 rate control structure. How and if that flexibility fits within the spectrum of IS missions needs to be considered. A more thorough investigation of RM2 capabilities as well as implementation assessments should be completed.

## REFERENCES

1. R. F. Rice, "Some Practical Universal Noiseless Coding Techniques," JPL Publication 79-22, Jet Propulsion Laboratory, Pasadena, CA., March 15, 1979.
2. R. F. Rice, "Practical Universal Noiseless Coding," SPIE Symposium Proceedings, Vol. 207, San Diego, CA., Aug. 1979.
3. R. F. Rice, A. P. Schlutsmeyer, "Software for Universal Noiseless Coding," Proceedings of 1981 International Conference on Communications, Denver, Colorado, June 1981.
4. R. F. Rice, A. P. Schlutsmeyer, "Data Compression for NOAA Weather Satellite Systems," Vol. 249, Proceedings 1980 SPIE Symposium, San Diego, CA., July 1980.
5. R. F. Rice, "An Advanced Imaging Communication System for Planetary Exploration," Vol. 66, SPIE Seminar Proceedings, Aug. 21-22, 1975, pp. 70-89.
6. R. F. Rice, "A Concept for Dynamic Control of RPV Information System Parameters," Proceedings 1978 Military Electronics Exposition, Anaheim, CA., Nov. 1978.
7. R. F. Rice et. al., "Block Adaptive Rate Controlled Image Data Compression," Proceedings of 1979 National Telecommunications Conf., Washington, D.C., Nov. 1979.
8. E. E. Hilbert, "Joint Pattern Recognition/Data Compression Concept for ERTS Multi-Spectral Imaging," Proceedings SPIE Seminar, Vol. 66, pp. 122-137, Aug. 1975.
9. R. F. Rice, "Channel Coding and Data Compression System Considerations for Efficient Communication of Planetary Imaging Data," Chapter 4, Technical Memorandum 33-695. Jet Propulsion Laboratory, Pasadena, CA., June 15, 1974.
10. R. F. Rice, "Potential End-to-End Imaging Information Rate Advantages of Various Alternative Communication Systems," JPL Publication 78-52. Jet Propulsion Laboratory, June 15, 1978.

11. K. Y. Liu, J. J. Lee, "An Experimental Study of the Concatenated Reed-Solomon/Viterbi Channel Coding System Performance and its Impact on Space Communications," JPL Publication 81-58. Jet Propulsion Laboratory, Aug. 15, 1981.
12. P. Ready, P. Wintz, "Information Extraction, SNR Improvement and Data Compression in Multispectral Imagery," IEEE Trans. Commun., Vol. COM-21, No. 10, Oct. 1973.