HYPERSPECTRAL IMAGE COMPRESSION

Dr. William H. Hallidy Jr. and Michael Doerr
Systems & Processes Engineering Corporation

ABSTRACT

Systems & Processes Engineering Corporation (SPEC) compared compression and decompression algorithms and developed optimal forms of lossless and lossy compression for hyperspectral data. We examined the relationship between compression-induced distortion and additive noise, determined the effect of errors on the compressed data, and showed that the data could separate targets from clutter after more than 50:1 compression.

KEY WORDS

Compression, Rice, Wavelet, Hyperspectral, Target/Clutter, Bit Error

INTRODUCTION

Hyperspectral imaging sensors and imaging spectrometers offer significant utility for a wide range of commercial and defense applications in the remote sensing area. These instruments image a scene over a large number of discrete, contiguous spectral bands such that a complete reflectance spectrum can be obtained for the region imaged. For example, the Hyperspectral Digital Imagery Collection Experiment (HYDICE), a high resolution, airborne imaging spectrometer, collects 210 contiguous spectral components over a range of 400 - 2500 nm for each of 320 separate spatial locations in a scene. Such a system generates a prodigious amount of information in a relatively short time. Assuming that the sensor array is digitized at 12 bits/pixel, with a frame time of 17.3 ms, we see that data is generated at over 5.8 Mbytes/s. An hour of continuous collection will provide a spectral/spatial hypercube of over 20 Gbytes of data. Hyperspectral instruments planned for future missions, such as Warfighter, HRST, ASRP and MARS, will significantly increase these numbers. Thus, it is challenging to store or transmit the data. Data compression can alleviate this problem. Data compression will also provide extended mission time over target capability due to the more effective use of on-board data storage and telemetry.

Systems & Processes Engineering Corporation (SPEC) has applied its experience with compression techniques to hyperspectral image (HI) data. In this report, we present some

of the findings of our research. We compared a number of compression algorithms and developed forms of both lossless and lossy compression for hyperspectral data. We also examined the relationship between compression-induced distortion and additive Gaussian noise, determined the effect of bit errors on the compressed HI data and showed that the data could separate targets from clutter after more than 50:1 compression.

SPEC applied several data compression techniques to HI data to establish optimal forms of both lossless and lossy compression. For lossy compression, a study of the effect of quantization distortion on the statistics of the data and a comparison of this distortion to the effects of noise was made. Quantitative studies were made of the relationship between the rms error due to distortion and the number of eigenvalues that were available for target detection in clutter. The study also included a determination of the effect of bit errors on the compressed images. Both lossless and lossy compression techniques were applied to samples of the HI data cube. The feasibility study successfully demonstrated data compression with a test set of AVIRIS HI data.

HYPERSPECTRAL IMAGING

Modern hyperspectral imaging systems employ a two-dimensional detector array. One dimension provides a linear array of spatial image pixels, while the other records hyperspectral data for each pixel. The spectrum is created by means of dispersive elements that allow acquisition of spectral data over hundreds of discrete spectral bands. The raw data for a single spatial pixel has the form of a continuous spectrum. As the camera platform moves, the camera scans the ground in a push-broom fashion to produce a data cube of hyperspectral images for on-board storage or immediate telemetry.

TARGET DETECTION IN CLUTTER

Hyperspectral Imagery usually contains objects of interest, ‘targets’, superimposed on a random background called ground clutter. In order to separate the target from the clutter in the scene, a statistical technique is frequently employed that involves using the covariance matrix of the data. The covariance matrix, **C**, is defined from the image. The matrix is symmetric and positive semidefinite. In practice, we expect **C** to be positive definite in all but the most exceptional instances. Thus, **C** may be diagonalized by a similarity transform with an orthonormal matrix. The diagonalized values are just the eigenvalues of **C**. The columns of the orthonormal matrix are its eigenvectors. For a symmetric, positive semidefinite matrix, the singular value transform and the eigenvalue decomposition of the matrix can be made identical. We use the singular value decomposition in our analyses.

The eigenvectors of the covariance matrix represent the directions in which the spectral scatter data is completely decorrelated. They are linear combinations of the basis defined

by the spectral bands. The eigenvalues define the variances in the eigenvector directions, so that for N_r spectral bands, a hyperellipsoidal error figure defines the standard deviation of the data in any direction. The ellipsoidal shape of the error figure can make it possible to separate a target from its clutter even when it is concealed in clutter in each of the spectral bands employed.

LOSSLESS COMPRESSION ALGORITHM

After comparing several lossless algorithms, we decided to use the Rice algorithm for lossless compression and, with some modification, as the lossless backend to our lossy compression algorithm. The Rice algorithm was designed to compress image data from space. It works in a raster fashion, organizing the image lines into blocks of pixels, usually with 16 pixels per block. Its results are comparable to those that would be achieved by designing a Huffman tree for each member of a set of entropy ranges and applying the tree that best fit the entropy for each block of pixels. However, the Rice code does not require such a set of trees.

One implementation requires a reference pixel at the beginning of each row in the image. Each pixel on the row thereafter is differenced with the preceding pixel to reduce redundancy. The difference values are mapped to a set of positive integers called sigma values. These sigma values create a probability distribution that approximates the Laplacian distribution. A comma code is formed from the sigma values. This code is concatenated for each block of pixels to form what is called a Fundamental Sequence (FS). The code, now renamed Universal Source Encoder for Space (USES), has been greatly improved from its original form by Dr. Pen-Shu Yeh of NASA Goddard.

The sigma values for a block are stored in a FIFO buffer while the bit counter determines which of 15 options will produce the shortest code. This is possible because the code length may easily be determined from the sigma values. The winning option code is then concatenated with appropriate identification bits and sent on.

LOSSLESS COMPRESSION RESULTS

In Table 1, we show the results obtained by lossless compression of the AVIRIS Moffett Field data. The Rice algorithm was used to obtain the tabulated results. In applying the algorithm, we chose to require eight pixels per block. This entailed the addition of two extra columns of zeros to the original frame. The average compression obtained over the data was 2.20:1.

Frame	50	100	150	200	250	300	350	400	450	500
Compression ratio	2.11	2.10	2.30	2.30	2.30	2.30	2.29	2.17	2.15	2.06

Table 1: Compression ratios of frames of Moffett Field scene 3 obtained with Rice

LOSSY COMPRESSION ALGORITHM

All lossy compression algorithms achieve their large compression ratios through some form of quantization. This is quantization that occurs in addition to the initial quantization that takes place in the A/D converter. For example, A/D conversion maps pixel values from an intractably large number of states (zero to full well of electrons) to a much smaller number of states (e.g., 0-256 for an 8-bit A/DC).

SPEC has compared the performance of lossy JPEG and Wavelet Compression techniques. We compressed frames of AVIRIS data to approximately the same degree using a JPEG Discrete Cosine Transform and using the wavelet transform. The final stage of lossless compression was performed with SPEC's modification of the Rice algorithm for each algorithm. The resulting compression ratios and relative errors are shown in Table 2. The results shown led us to choose the Wavelet Transform for compression of the Hyperspectral Image data. We also conducted a study of the choice of wavelet basis that resulted in our deciding to use the Daubechies 5 wavelet for the lossy compression algorithm

Algorithm	Compression Ratio	Relative Error (%)
JPEG	50.4	0.46
Wavelet	58.6	0.36

Table 2: Comparison of lossy algorithms

LOSSY COMPRESSION RESULTS

We sampled the compression on frames of the entire AVIRIS data cube with the Daubechies 5 basis and uniform 8-bit quantization. The results of this sampling are shown in Table 3. The mean compression ratio over the samples is 59.8:1 with a mean relative rms error of 0.43%. The median values of the compression ratio and rms error are 60.3:1 and 0.38%, respectively. From the table, it appeared that frame 300 of the data cube gave results that were fairly representative of the average over the samples. We therefore used it as an exemplar of the cube's compression behavior.

Frame	50	100	150	200	250	300	350	400	450	500
Compression ratio	87.6	46.0	60.7	64.1	65.3	58.6	62.8	58.1	59.9	50.3
Relative rms error (%)	0.75	0.41	0.37	0.36	0.38	0.36	0.37	0.43	0.38	0.46

Table 3: Sampled compression of Moffett Field scene 3 data

We compressed AVIRIS data from frame 300 of the data cube employing a wavelet transform with the Daubechies 5 basis. The final lossless compression was performed with a SPEC modified form of Rice compression. Uniform bit quantization was used in each case on the transformed coefficients at each of the 4 levels of decomposition employed in the wavelet transform. Table 4 shows the compression results for the data, for uniform quantization levels from 12 to 6 bits. Non-uniform sampling can be used to trade off some compression for reduced rms error.

Qbits	12	11	10	9	8	7	6
Compression ratio	10.8	15.6	23.0	35.4	58.6	91.7	146.0
rms error (%)	0.04	0.07	0.15	0.23	0.36	0.64	0.90

Table 4: Moffett Field compression data

MEASURE OF DISTORTION

There is still a question as to what level of distortion is acceptable. Smaller distortion values should mean more reliable data, but what constitutes a sufficiently small value for HI data that is to be used for target detection and terrain classification? In the following section, we suggest that the answer to the question may be found by examining the effects of compression on the eigenvalues of the covariance matrix. In order to develop a valid measure of distortion based on the eigenvalue decomposition of the covariance matrix, we compared of the effects of compression induced distortion on the eigenvalues with effects induced by the addition of white Gaussian noise.

COMPRESSION EFFECTS ON EIGENVALUES

We recall that the eigenvalues of the covariance matrix give the variance of the scatter plots in the directions of the corresponding eigenvectors. Target detection in clutter depends on the presence of target points outside of the error figure of the clutter's scatter plot. Figure 1 compares the 50 largest eigenvalues obtained after compression with those of the original image. Since eigenvalues extend over a very large dynamic range, we have

employed a semi-logarithmic scale to better show their variation. The original values are indicated as black + signs. The compressed data eigenvalues are shown for uniform quantization at levels from 6 to 12 bits per coefficient, as depicted in the accompanying legend. As may be seen from this figure, at a point that depends on the level of quantization, the eigenvalues of the compressed data can grow an order of magnitude or more larger than those of the original data. On the other hand, the smallest eigenvalues of the compressed data were found to fall about two orders of magnitude below those of the original data. We also compared individual eigenvalues of compressed data with those of the original data. In all cases, we found that the largest eigenvalues for compressed data were in good agreement with the original data.

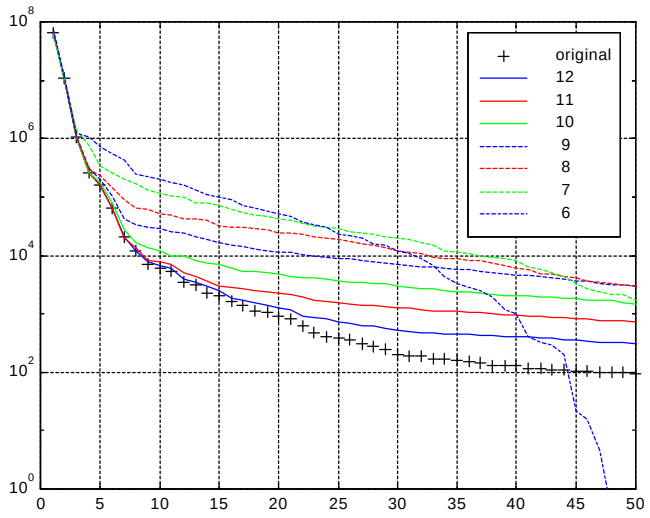

Figure 1: Compression effects on eigenvalues of AVIRIS data

Such discrepancies in agreement may have consequences for target detection. We shall assume that the eigenvalues and eigendirections of uncompressed data for the clutter environment are used to form an error figure that will be used for target detection decision making.

If compressed data is used to send real time information about possible targets in the environment, changes in the eigenvalues or in the eigendirections can adversely affect the results. Larger eigenvalues correspond to greater scatter of the clutter. This could cause some clutter pixels to move outside of the error figure, leading to false positive target detection results. Smaller eigenvalues, indicating less variance in the scattered data, might cause target pixels to move within the clutter error figure, resulting in a target detection failure. Rotation of the eigendirections of the compressed data with respect to the original data can be shown to lead to similar conclusions. Therefore only the eigenvalues and eigenvectors with the least amount of change will be useful in target detection and terrain classification.

NOISE EFFECTS ON EIGENVALUES

We have tried to assess the significance of the distortion described in terms of the relative rms errors in the compressed data by making a comparison of the eigenvalues of the data with those produced by the addition of white Gaussian noise to the original data. In order to facilitate the comparison, we made the standard deviation of the noise equal to the value of the distortion for each quantization level compared.

Figure 2 illustrates the results of this comparison for quantization levels 6, 8, 10, and 12. The original data eigenvalues are again displayed with black + signs. The legend accompanying the figure explains the labels used to display the eigenvalues of the remaining parameters. The parameters are labeled according to their quantization level for the compressed data. The noise data parameters are indicated by an appended n. The quantization level indicated corresponds to a distortion equal to the standard deviation chosen. The distortion values are listed in row three of Table 3.

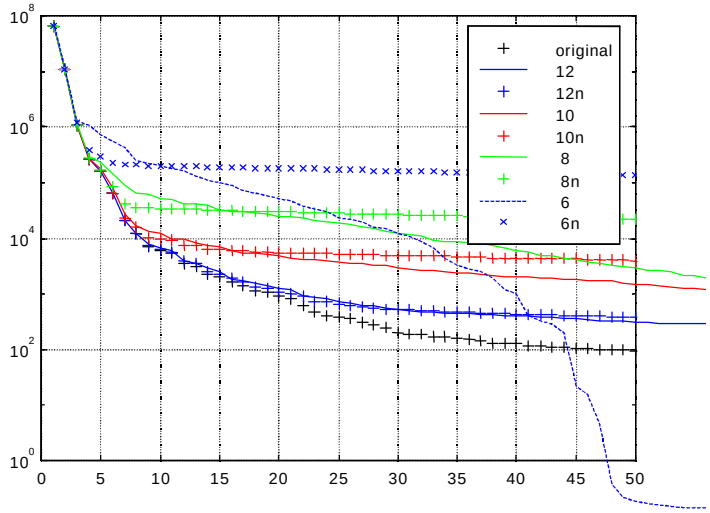

Figure 2: Compressed versus noisy eigenvalues for AVIRIS data

As shown in Figure 2, the eigenvalues of the noisy data depart from those of the original data at about the same location as do those of the corresponding compressed data. However, the slopes of the noisy data values flatten out thereafter. This corresponds to eigenvalues and eigendirections that are dominated by noise. The effective number of spectral degrees of freedom is actually equal to the number of significant eigenvalues. Flattening out of the eigenvalues makes the error figure nearly spherical. Under such circumstances, for a target to be discriminated from clutter dominated by noise, its spectral value would have to dominate the clutter in at least one of the available spectral bands.

It may therefore be argued for the compressed data case that the number of unaffected spectral values determines the effective number of spectral degrees of freedom. As the

eigendirections are linear combinations of the spectral bands, we see that the effective number of independent spectral bands equals the number of eigendirections and eigenvalues that are significant. For the Moffett Field data that we analyzed, we found that the number of significant eigenvalues varied roughly linearly with a "signal to noise ratio" defined as:

$$SNR = 20 \cdot \log_{10}(s)$$

where σ is the compression-produced distortion. More data would have to be examined to determine whether a quantitative relationship between the SNR and the number of eigenvalues holds in general.

TRANSMISSION BIT ERROR ANALYSIS

We simulated the effect of uncorrected bit errors arising in transmission of the data on the wavelet-transformed coefficients of AVIRIS data for different assumed Gaussian distributed error rates. We also determined single bit worst case errors for each level of the decomposition by flipping a most significant bit in one of the level's coefficients.

Figure 3a shows an original spatial/spectral image of AVIRIS, Moffett Field data. Figure 3b shows the same image affected by random errors with a probability of error, p_{err}, of 10^{-6} per coefficient. In this image there is a bit error in each of three coefficients, although only one error is large enough to be visible. The spread in the error is due to the filtering and decimation actions of the transform. For the worst case errors, the spread in the image depends on which of the coefficients is affected.

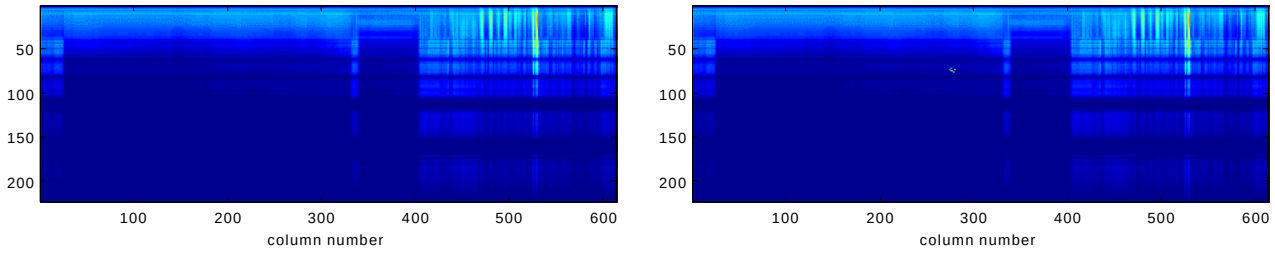

(a) (b)

Figure 3: (a) Original Spatial/Spectral Image (b) Image with single bit errors in three coefficients

We also determined the effect on the data of random bit errors where the assumed probability of error was 10^{-5} and 10^{-4} per pixel. These values introduced on the order of 10 and 100 bit errors in the coefficients, respectively. Based on these results, it appears reasonable to suggest that the image frames would still be useable for uncorrected bit errors in the compressed coefficients at a level of up to 10^{-5} errors/coefficient.

The spreading of bit errors in the image depends on the coefficients that are affected. Figure 4 shows a matrix arrangement of our level four, wavelet decomposition. The number of coefficients in each subdivision is approximately proportional to the area of the subdivision. Because of interpolation, the errors in the higher numbered levels spread more in the image, but errors at all levels also spread because of the filtering action of the wavelets.

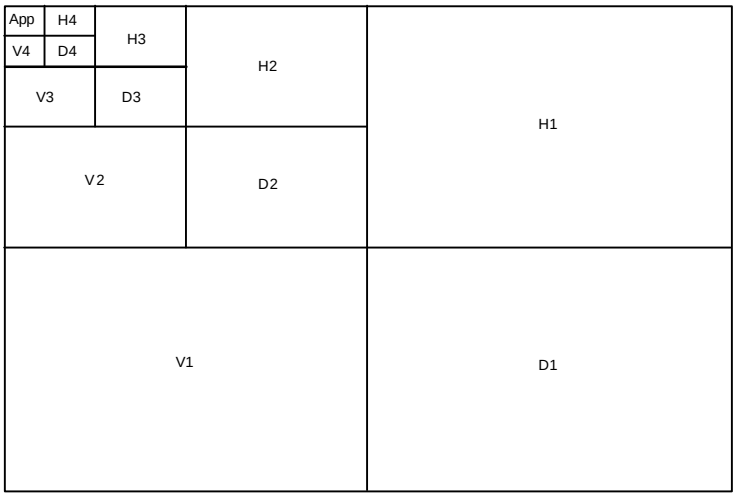

Figure 4: Two Dimensional Arrangement of a Level four decomposition

We examined the effect of worst case errors at various levels as applied to AVIRIS data. A wavelet transform was applied to the frame without compression. A most significant bit (msb) error was applied to one of the coefficients and the values were inverse transformed to produce the resultant frame.

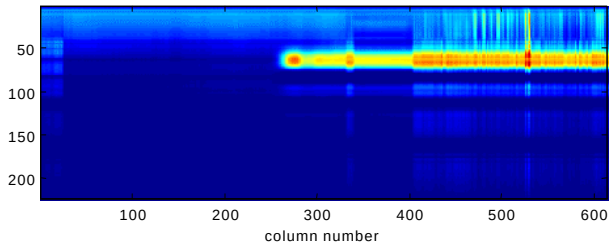

Figure 5: Error in msb of one approximation coefficient

Comparing Figure 3a with Figure 5 shows the effect of a single msb error in one of the approximation coefficients. This is truly the worst possible single bit error, because the approximation coefficients are differenced rowwise, thus allowing the error to propagate across the remainder of the row. This effect could be ameliorated or even eliminated by allowing more reference pixels per row or by eliminating differencing altogether in applying the lossless backend compression algorithm. Some loss of compression would accompany this however. The spread across bands shown in Figure 5 is caused by the

filtering action of the wavelet inverse transform combined with the effect of multiple interpolation operations.

The degree of spreading of a single bit error depends on the length of the wavelet filter coefficients. For the wavelet basis we used, an error in one of the level 1 coefficients appeared to be confined to the nearest pixels in each direction in the image. The number of pixels in each direction influenced by the error would approximately double for an error in each succeeding level's coefficients.

The probability of a random bit error occurring in a particular subdivision is proportional to the size of the subdivision. From Figure 4 it is clear that the probability of an error occurring in a coefficient in the level 4 approximation subdivision is about 1/4 that of its occurring anywhere in level 4. The probability of an error anywhere in level 4 is about 1/4 that of an error in levels 3 or 4, etc. Therefore, the probability of an error occurring in one of the approximation coefficients is about $1/256^{th}$ that of an error occurring anywhere in the coefficients.

TARGET/CLUTTER SEPARATION

In order to determine the effect of compression on target/clutter separation, we chose a row that had been identified as containing a target as our target/clutter data. We used nearby rows that were devoid of targets to define the background clutter data. We created a 3D-scatter plot of the clutter data for three selected spectral bands. The axes of the error ellipsoids were determined and compared with the axes for a 3D-scatter plot of the target/clutter data. A small amount of rotation was present between the two scatter plots in two of the eigendirections. The third direction appeared to be coincident between plots. By identifying the scatter points contributed by the target, we found that most of the target points lay outside the 3-sigma error ellipse defined by the clutter data. We also found that all the target points lay within the error ellipse projected onto two of the three spectral bands. The image was then compressed by about 50:1 and the resulting scatter plot was again examined. We found very little difference between the compressed and uncompressed eigendirections. Slight changes in the eigenvalues were observed. Most of the target points remained outside of the error ellipse defined by the clutter. So, in this instance, at least, it was possible to separate the target from the clutter after compressing the image by about 50:1.

CONCLUSIONS

Our investigations described in the preceding sections led us to conclude that Rice/USES was the best algorithm to use for lossless compression of hyperspectral data. Rice has been adopted as a CCSDS standard. It is non-bursty, relatively, simple to produce in hardware,

and robust against bit error when packetized. With a compression ratio of somewhat greater than 2:1, it can permit significant reduction in storage requirements, or, for the same amount of storage, it can extend the time over target by a factor of two with no loss in image quality. We have also created a modification of the Rice algorithm for use as a lossless backend entropy coder for our lossy algorithm.

Our study of lossy compression techniques led us to choose a four level, Daubechies 5 basis, wavelet transform with a SPEC modified Rice entropy coder backend, as a lossy algorithm that was ideally suited to compression of hyperspectral image data. With this algorithm we were able to achieve almost 60:1 compression with less than 0.4% distortion of AVIRIS Moffett Field data. This amount of compression would permit BPSK telemetry of the data over a 5 MHz bandwidth channel.

We also made an extensive analysis of the effect of distortion on the data. This included a comparison of the effects of compression distortion with those of noise. The effect of bit errors, including worst case errors was also analyzed. In addition, data was used to show that after compression it was still possible to discriminate targets from clutter.

SELECTED RELEVANT LITERATURE

1. Rice, R.F., "Practical universal noiseless coding", SPIE Vol. 207, Application of Digital Image Processing III, pp. 247-267, 1979.
2. Yeh, P.-S. and Miller, W.H., "Application Guide for Universal Source Encoding for Space", NASA Technical Paper 3441, December, 1993.
3. *Universal Source Coding for Data Compression* Draft Recommendation for Space Data Systems Standards, CCSDS 121.0-R-1, Red Book, Issue 1, Washington, D.C.: CCSDS, November 1995.
4. Daubechies, I., The Wavelet Transform, Time-Frequency Localization and Signal Analysis, IEEE Trans. Information Theory, 36(5),961(1990).
5. Mallat, S.G., "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation", IEEE Trans. Pattern Anal. Machine Intelligence, PAMI 11(7), 674, (1989).
6. Schaum, A. and Stocker, A., "Subclutter target detection using sequences of thermal infrared multispectral imagery", SPIE, 3071, 12-22, (1997).
7. Landgrebe, D., "Multispectral Data Analysis: A Signal Theory Perspective", available from http://dynamo.ecn.purdue.edu/~biehl/MultiSpec/documentation.html, pp 46, (1998), other related papers by this author also available at the above site.
8. Theiler, J., "Finding Weak Signatures in Cluttered Backgrounds Using Covariance Matrix Methods", available from http://sst10a.lanl.gov:80/~jt/Preprints/#findsig, pp. 11, (1997).

BIOGRAPHY

Dr. William Hallidy, holds a Ph.D. in Theoretical Physics from the University of Pittsburgh. He has more than 25 years of diverse research and system development experience with prominent government, industrial, and academic institutions such as the Naval Weapon Center, ERIM, and Georgia Tech. He is employed by Systems & Processes Engineering Corporation, where he now holds the position of Chief Analyst.

At SPEC, Dr. Hallidy is responsible for system analysis, modeling and algorithm development. He has analyzed compression, forward error correction and modulation techniques for a digital video telemetry module, having completed an extensive investigation of lossless and lossy algorithms for data and image compression, and has applied these techniques to the compression of hyperspectral interferometric imagery. Dr. Hallidy's image analysis laboratory supports the quantitative and qualitative analysis of digital imagery for pre- and post-compression evaluation. Dr. Hallidy has also developed and utilized additional models for communications link budget predictions and signal-to-noise analysis for both communications and sensor systems and has investigated a number of signal processing models for automatic target recognition.

He has developed a new chemometric algorithm for the analysis and extraction of chemical components from mixtures with phase resolved fluorescence data and has modeled the signal to noise performance of a complete remote sensing system for fluorescence spectroscopy and contributed to the design of its optical system. Dr. Hallidy has also simulated the performance of a digital RF memory system and has explored different analytical methods for predicting the electromagnetic and thermal characteristics of very high speed digital integrated circuits. He has developed moment method based software in C to predict the interconnect capacitance and inductance of integrated circuits and has investigated slope models for gate level timing analysis of GaAs based digital MESFET designs.

Dr. Hallidy has participated in the development of analytical models for synthetic aperture radar (SAR), including bistatic spotlight SAR, and for synthetic aperture imaging in the IR and visible wavelengths and has modeled the effects of dispersion on the vertical atmospheric propagation of chirped laser pulses. He has also modeled plasma induced phase distortion effects on two-way signals and scattering by rough surfaces and complex targets in the RF, IR and visible frequencies. He has experience in the design of both analog and digital circuits.

For many years he has been active in the area of electromagnetic propagation, scattering and imaging of RF, IR and optical phenomena and has authored numerous reports and technical papers on these subjects.