SAND97-1928C
SAND--97-1928C
CONF-980524--
RECEIVED
JAN 29 1998
OSTI

INTELLIGENT LOW RATE COMPRESSION OF SPECKLED SAR IMAGERY

IEEE 1998 NATIONAL RADAR CONFERENCE*

Robert W. Ives[1], Paul Eichel[1] and Neeraj Magotra[2]
[1]Sandia National Laboratories, Albuquerque, New Mexico, USA 87185
[2]EECE Department, University of New Mexico, Albuquerque, New Mexico, USA 87131

ABSTRACT

This paper describes a compression technique under development at Sandia National Laboratories for the compression of complex synthetic aperture radar (SAR) imagery at very low overall bit rates. The methods involved combine several elements of existing and new lossy and lossless compression schemes in order to achieve an overall compression ratio of large SAR scenes of at least 50:1, while maintaining reasonable image quality. It is assumed that the end user will be primarily interested in specific regions of interest within the image (called "chips"), but that the context in which these chips appear within the entire scene is also of importance to an image analyst. The term "intelligent" is used to signify an external cuer which locates the chips of interest.

1.0 INTRODUCTION

The immediate goal of this research is to address the problem of transmitting massive amounts of high resolution complex SAR data from a remote airborne sensor to a ground station for exploitation by an automatic target recognition (ATR) system, in a real time environment. The system under development is for the *Intelligent Bandwidth Compression* (IBC) program funded by the Defense Advanced Research Projects Agency (DARPA). The constraints imposed by the problem under consideration demand a background compression on the order of 100:1, chip compression on the order of 7:1, and a chip rate not to exceed approximately 67 chips per million pixels. The ATR system results might then be presented to an image analyst who, using the contextual information from the entire SAR image, makes final target determination. At the ratios desired, the compression can be achieved by applying one of several existing lossy techniques to the image as a whole: however, at such high ratios (and possibly even at low compression ratios), an intolerable amount of vital information can be lost and/or unacceptable coding artifacts introduced. These two effects may severely degrade the ATR system's performance, not to mention the perceptual quality of the image from the analyst's perspective.

The term "intelligent" is utilized to signify the presence of an external cuer which designates regions of interest (termed "chips") within the SAR scene. The use of this cuer leads us to divide the image into two segments, "background" and "target", each of which are subjected to different compression algorithms. The use of multiple compression schemes within the same image we refer to as "multi-modal" compression. "Target" can mean either a square chip or a group of pixels within a chip, while "background" is simply the entire image as a whole. "Targets" are compressed in one of two ways: either losslessly or with slight loss, while the background is significantly compressed. The use of a cuer with these two distinct paths allow high fidelity coding where it is needed, and very low rate coding where it is not needed, yielding low overall rate while maintaining high quality on regions of interest.

2.0 COMPRESSION: CHIP PROCESSING

Ideally, perfect fidelity of the decompressed chips is desired: this would cause no performance degradation on the downstream chip processing. For perfect fidelity, chips could individually be processed losslessly using perhaps one of the methods in [1], achieving a compression ratio on the order of 2:1. However,

*Presented at the IEEE 1998 National Radar Conference, Dallas, Tx, 12-13 May 1998

[1] Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

19980420 069

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

for the overall compression ratio required in this scenario, this is not sufficient. Therefore, chips *must* be compressed with some loss of information, i.e., either (1) the entire chip is compressed with minimal, acceptable loss ("transform domain quantization") or (2) only a certain percentage of the pixels on the chip are transmitted losslessly ("pixel segmentation"). These methods are described in more detail in [2]-[4], with brief descriptions given in the following two sections.

2.1 TRANSFORM DOMAIN QUANTIZATION

The primary method to compress the chip as a whole is based on the quantization of the coefficients after applying a transform. Transform domain quantization is lossy over the entire chip. It is done in such a way, however, to minimize the loss and maintain good chip fidelity [4]. This is done by taking advantage of two key aspects of SAR data, namely oversampling and the *much* reduced effects of quantizing in the transform domain compared to the image domain. And this method is unlike the common application of a transform, where the goal is to concentrate the energy of the input image into a small number of coefficients, allowing smaller valued coefficients to be severely quantized or zeroed out. In this method, the coefficients are quantized so as to limit the effects of harmonic component distortions, precluding ringing or blocking effects commonly seen when using the DCT [4]. Processing chips in this manner, the data to be transmitted includes (1) the compressed background, (2) the transform domain data and (3) the overhead data. This method generally achieves a lower overall compression ratio than using pixel segmentation.

In its most basic form, this method is the two-dimensional fast Fourier transform (FFT) of the complex input chip which is nominally a power of two in size in each dimension. But due to the oversampling performed on SAR imagery in both the range and azimuth directions, only a fraction of the FFT coefficients are required to represent the image domain chip: in fact, the size of the chip can be reduced by a factor equal to the product of the oversampling ratios in the range and azimuth dimensions with only slight loss. These ratios generally fall into a range of values of 1.1-1.6, and are equal to 1.333 for the imagery in the IBC program, which would in itself allow data reduction of a chip by 1.78. At this point, the oversampling redundancy is removed by extracting a sub-chip from this data, the size dependent on the oversampling ratios. Note however, that by using the FFT, the values used to represent the chip would change from four bytes per pixel (I/Q, signed short) to eight bytes per pixel (floating point complex), which would not provide compression unless quantized. The second element of this technique, then, is quantization of the resulting coefficients. This step will produce some loss in the reconstructed chip. The dynamic range of this FFT data is reduced by dividing by a scale factor related to the variance (this scale factor is transmitted for each chip), and the result is quantized to eight bits per pixel (four signed bits in-phase/four signed bits quadrature) with a uniform quantizer. This quantizing introduces an additional 4:1 compression in each chip. By further coding of the quantized coefficients with an arithmetic coder, a further gain of approximately 8% can be realized. As an example, for a chip from a 32-bit per pixel complex IBC image whose oversample ratios are 1.33 and whose coefficients are quantized to 8 bits per pixel the chip can be compressed at $1.33*1.33*(32/8)*1.08=7.64$ with minimal loss.

2.2 PIXEL SEGMENTATION

For this method, two problems must be addressed: determining which pixels within a chip should be labeled as belonging to the target of interest, and the number of pixels that should be preserved. Given an external cuer, specific chips have been designated as possible targets in the SAR image and it is the man-made objects (specular targets) which generally give the largest magnitude returns to the radar. Accordingly, for an ATR system which attempts to locate and identify man-made objects (such as tanks, missiles or other vehicles), collections of large magnitude pixels within this chip would tend to belong to the possible target. A simplistic approach to locating target pixels would then be to find the largest magnitude pixels, the number of pixels which may be designated as target pixels dependent on the number of chips and the required compression ratio. These pixel values will be transmitted losslessy, along with a compressed bit map of the entire image to keep track of their locations. Note that such a bit map, with small patches of ones in chip locations and zeros over the vast majority is well suited for run-length coding. Simply using the largest magnitude pixels, however does not take advantage of the correlation which exists between pixels on a target, and may call out large-valued clutter as target pixels, wasting bandwidth. A better method of pixel segmentation would measure the correlation between pixels on the chip, then preserve those pixel values corresponding to the highest correlation. This has been achieved with the two-dimensional adaptive correlation enhancer (2DACE) of [2].

The 2DACE provides a recursive estimate of spatial correlation within a chip. For each pixel within the input chip, the correlation coefficients are computed within a local neighborhood. The output of the 2DACE is a floating point map representing the computed normalized correlation coefficients, with higher values representing more correlation among neighboring pixels. From this correlation map, those locations with the largest correlation are chosen as target pixels on the SAR image chip. The number of pixels which can be preserved is necessarily based on the number of chips present in a particular image and the bandwidth, although the relation is not linear, since the bit map also changes shape as the number of pixels per chip changes. Therefore, the data to be transmitted for any input image is (1) the compressed background, (2) the compressed bit map, (3) the target pixel values, and (4) overhead data, including the chip locations, image size, etc.

3.0 BACKGROUND PROCESSING

A detailed description of the background compression algorithm is given in [5], and will be briefly described here. Because SARs are *coherent* imaging systems, the images they produce are speckled. While the speckle noise generally conveys no real information, being a high entropy process it severely degrades the compression/ decompression results. The compression artifacts prominent in highly compressed SAR imagery are directly related to the high frequency speckle inherent in these types of images. If the speckle is removed/reduced prior to compression, the artifacts in the decompressed image will be reduced. In [5], we derive a maximum likelihood estimator for the mean backscatter coefficient function. It is this speckle "mean" image that is compressed, which preserves the statistics of the speckle such that the decompressed image will have approximately equal local speckle statistics to the original. Since the speckle mean function (image) is compressed vice the original image, the high frequency noise is substantially reduced prior to the compression engine resulting in minimized coding artifacts in the decompressed image. For the IBC system, the speckle mean background image is subjected to a fixed 64:1 compression.

Shapiro's Embedded Zerotree Wavelet algorithm [6] is used to compress the speckle "mean" image. We use the length 2/6 biorthogonal filter bank which was evaluated for image compression in [7]. It was found that many short filters (such as the Daubechy orthogonal 4-tap and 6-tap, the Haar 2-tap, and the biorthogonal length 5/3 and 2/6) outperformed longer filters in terms of mean-squared error (MSE) of the decompressed speckle "mean" image and generally provided approximately the same amout of error as each other: but in some instances the length 2/6 filter yielded an order of magnitude less MSE. As a modification to the basic algorithm in order to improve run time, in the first decomposition the outtermost three subands are not computed. This allows more bits to be allotted to the compression of the lower frequency components and since much of the high-frequency energy is removed by prefiltering, this was found to have little or no perceptual degradation in quality compared to when the outtermost subands were computed and compressed.

The background image is wavelet decompressed (the outtermost subands in the final recomposition are zeros), and using the additional radar parameters necessary to characterize the speckle statistics, complex speckle is added. It is into this respeckled, decompressed background image that the target chips (using transform domain quantization) or pixel values (using pixel segmentation) are inserted. It is important to note that for the background, no attempt is made to preserve the phase, although the respeckled background image is complex: only the phase of the target chips or target pixels are accurate. Using pixel segmentation, the bit map is decompressed and raster scanned. Whenever a target pixel is indicated, the next value from the pixel value file is inserted in the modified background. Using the transform domain quantization, each chip's data is scaled, inverse transformed and inserted in its entirety into the background. Because the statistics of the speckle is preserved in the background, a significant benefit of this compression scheme is that the chip data insertion is seamless, which prevents distracting visual artifacts from appearing in the final product. Figure 1 displays a portion of a 1 meter resolution SAR image of the Pentagon collected by the Sandia SAR carried onboard the U.S. Department of Energy's Airborne Multisensor Pod System (AMPS) aircraft. On the left is the original; in the center is the result if the detected image is subjected to a 64:1 EZW compression without any pre- or post-processing; and on the right is the result from the IBC system, also 64:1 but applied to the speckle mean image and with respeckling.

4.0 RESULTS

The quality of a SAR image compressed/decompressed in a lossy manner can be readily tested using the ATR system for which it was designed, and monitoring performance degradation. Preliminary testing of chips compressed with transform domain quantization on a Sandia ATR system showed no degradation. However, the quality is highly subjective when the human visual system is used to judge. Various measures have been used to quantify the quality of lossy compressed imagery, including SNR, MSE, etc., but these are probably not a reasonable metric when the environment contains uncorrelated speckle throughout, particularly when the preservation of particular speckle values is unimportant (e.g., background). In the case at hand, the object is to determine the quality degradation between two SAR images which are of the same size and are registered, since one is just a compressed/decompressed version of the other. From these two images a correlation map may be computed. This map will have the same dimensions as the original: for each pixel location, the correlation coefficient is calculated based on the pixels in a small neighborhood about it (5x5 for example) using common pixels in each image. In this manner, the sample complex spatial correlation coefficient for a pixel at location *(x,y)* between an original and decompressed SAR image can be computed from [8]

$$c(x, y) = \frac{\left|\sum_i \sum_j f(i, j) \cdot g^*(i, j)\right|}{\sqrt{\sum_i \sum_j |f(i, j)|^2 \cdot \sum_i \sum_j |g(i, j)|^2}}, \qquad (1)$$

where the indices i and j represent pixels in a neighborhood about pixel *(x,y)*, and f and g are the original and decompressed images. This correlation is measured over a neighborhood about each pixel in the chip, and will fall in the range [0.0,1.0]. Perfect correlation for a particular pixel is indicated by a value of 1.0 and occurs if every complex value within the neighborhood of the original chip is equal to the corresponding pixel value in the decompressed chip: a value of 0.0 indicates that they are completely uncorrelated. Equation (1) assumes a zero mean in the neighborhood, which is generally true for complex SAR images with sufficient neighborhood size.

As an example of this compression technique, the Pentagon image was input into to the the IBC compression system with four "targets" identified, and corresponding 64x64 chips were processed with the transform domain technique. A standard chip size for this resolution in the IBC system would be 32x32: oversized chips were used for viewability. The original image and resulting decompressed IBC image is included as Figure 2. The overall compression ration with four oversized chips is 63.78:1. To demonstrate the quality of this test image, the corresponding complex correlation maps is also included as Figure 3, scaled from [0.0,1.0] to [0,255]. The average complex correlation value for the four chips is 0.89. The correlation map shows clearly the location of the target chips, however, their seamless integration precludes their identification in the finished product (on the right in Figure 2).

5.0 REFERENCES

[1] R. Ives, N. Magotra and G. Mandyam, "Lossless Compression of Synthetic Aperture Radar Images", In *1996 IEEE Int. Sym. on Circuits and Systems,* Atlanta, GA, pp. 441-444, May 1996.

[2] K. Knittle and N.Magotra, "Two Dimensional Adaptive Enhancement and Detection of Linear Features in Noisy Images", *25th Asilomar Conference on Signals, Systems and Computers"* Asilomar, CA, Nov. 1991.

[3] R. Ives, P. Eichel and N. Magotra, "Application of Pixel Segmentation to the Low Rate Compression of Complex SAR Imagery", submitted to *IEEE International Geoscience and Remote Sensing Symposium '98*

(IGARSS '98).

[4] P. Eichel and R. Ives, "Compression of Complex Valued SAR Images", submitted for publication to the *IEEE Transactions on Image Processing.*

[5] P. Eichel and R. Ives, "Very Low Rate Compression of Speckled SAR Imagery", to be presented at *SPIE's 12th Annual Symposium on Aerospace/Defense Sensing, Simulation and Controls (AEROSENSE 98).*

[6] Shapiro, J. M. "Embedded image coding using zerotrees of wavelet coefficients", *IEEE Transactions on Signal Processing,* Vol. 41, No. 12, pp. 3445-3462, Dec. 1993.

[7] Villasenor, J.D., Belzer, B. and Liao, J. "Wavelet filter evaluation for image compression", *IEEE Transactions on Image Processing,* Vol. 4, No. 8, pp. 1053-1060, Aug. 1995.

[8] Jakowatz, C. et al. *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach..* Boston, MA: Kluwer Academic Publishers, 1996.

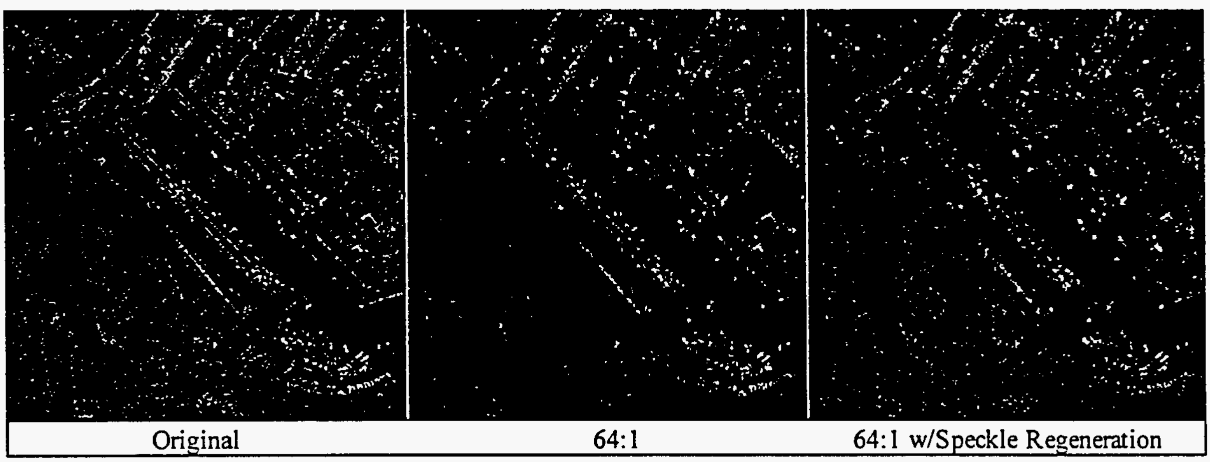

Figure 1: Effects of Pre-Filtering and Speckle Regeneration

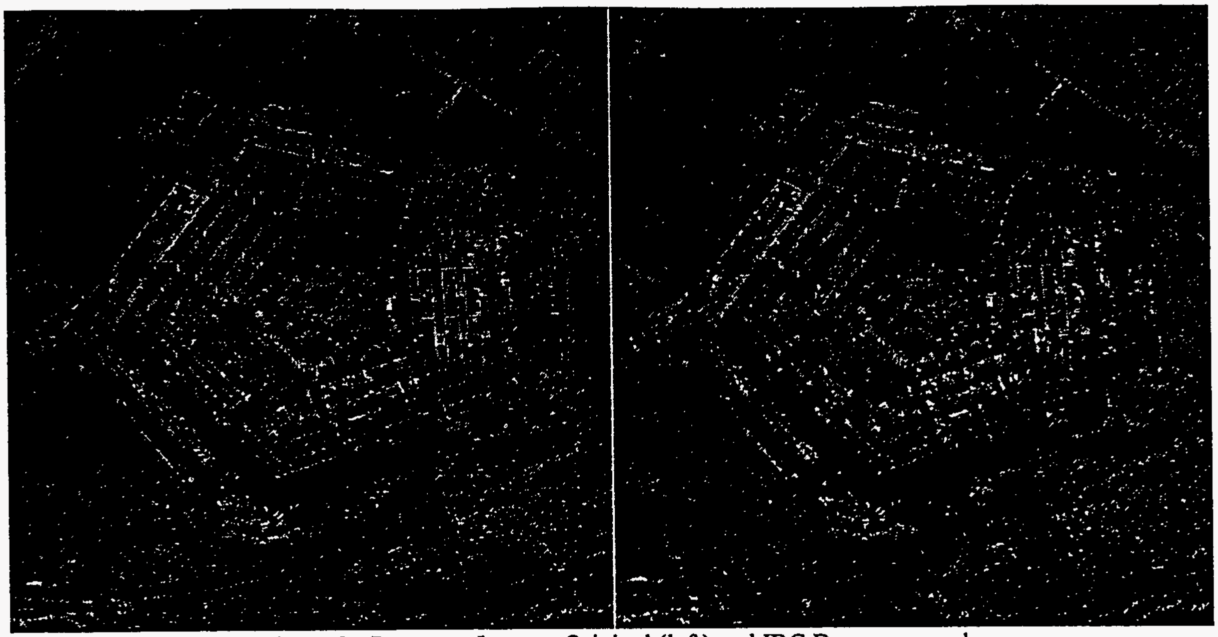

Figure 2: Pentagon Image: Original (left) and IBC Decompressed

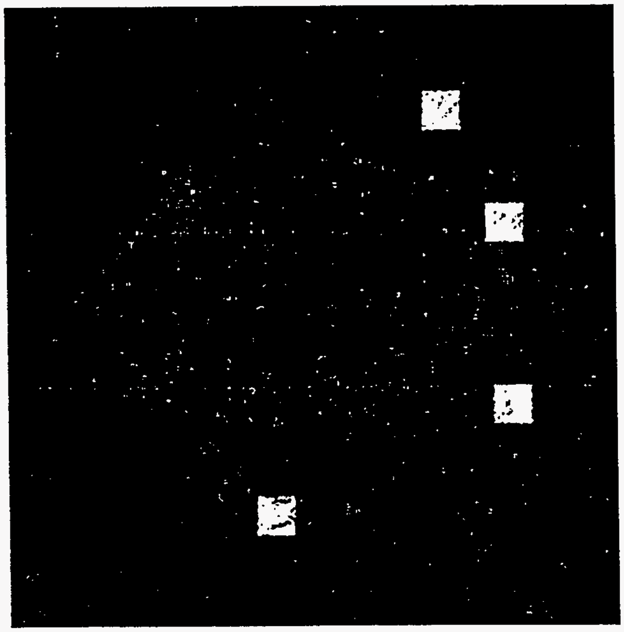

Figure 3: Complex Correlation of Pentagon Image

M98004097

Report Number (14) SAND--97-1928C
CONF-980524--

Publ. Date (11) 199805

Sponsor Code (18) DUS, XF

UC Category (19) UC-000, DOE/ER

DOE