14th European Signal Processing Conference (EUSIPCO 2006), Florence, Italy, September 4-8, 2006, copyright by EURASIP

NEAR-LOSSLESS DISTRIBUTED CODING OF HYPERSPECTRAL IMAGES USING A STATISTICAL MODEL FOR PROBABILITY ESTIMATION

Marco Grangetto, Enrico Magli, Gabriella Olmo

Dept. of Electronics, Politecnico di Torino
Corso Duca degli Abruzzi 24, 10129 Torino, Italy
e-mail: firstname.lastname@polito.it

ABSTRACT

In this paper we propose an algorithm for near-lossless compression of hyperspectral images based on distributed source coding (DSC). The encoding is based on syndrome coding of bit-planes of the quantized prediction error of each band, using the same information in the previous band as side information. The practical scheme employs an array of low-density parity-check codes.

Unlike other existing DSC techniques, the determination of the encoding rate for each data block is completely based on a statistical model, avoiding the need of inter-source communication, as well as of a feedback channel. Moreover, the statistical model allows to estimate the statistics of the currently decoded bit-plane also using the information about the previously decoded ones in the same band; this boosts the performance of the DSC scheme towards the capacity of the conditional entropy of the multilevel (as opposed to binary) source.

Experimental results have been worked out using AVIRIS data; a significant performance improvement is obtained with respect to existing DSC and classical techniques, although there is still a gap with respect to the theoretical coding bounds.

1. INTRODUCTION

Distributed source coding (DSC) has recently attracted a lot of interest in the signal processing community [1]. DSC refers to the problem of encoding two (or more) correlated sources, say X and Y, without letting them communicate with each other, i.e. avoiding joint encoding. Since the joint encoder requires explicit decorrelation of the sources, there exist several applications where it is undesirable to use one. For example, when the encoder computational resources are scarce, decorrelation of a complex data set such as a video sequence or a multichannel image can be a burdensome task. Moreover, the sources could be located in different places, as in sensor networks applications, making it very costly to transmit all the sources to a central joint encoder. However, it has been shown that, for lossless compression, separate encoding of correlated sources can be performed without any performance loss with respect to joint encoding, provided that *joint decoding* is carried out [2]; this combination of separate decoding and joint decoding is referred to as DSC. Similar results exist in the lossy case [3], stating that the lossy DSC problem may exhibit a performance loss with respect to the joint encoder; however, the loss is no larger than 0.5 bpp [4], and there exist no-loss cases, e.g. if the sources can be modeled as $X = Y + N$, with N a zero-mean Gaussian process independent from Y.

These results have been used to revert traditional coding paradigms employing a complex joint encoder, which does all the source modeling, and a light decoder. In DSC, there is little or ideally no encoder modeling, while this is done at the decoder side, allowing to obtain a scheme consisting of a light encoder and a complex decoder. Applications to lossy video coding [5] and lossless coding of hyperspectral images [6] have shown the feasibility of this approach.

One serious issue with many DSC schemes lies in the conditional entropy estimation stage. To clarify this, we employ the simplest lossless DSC scheme (also known as Slepian-Wolf coding), where Y is encoded at rate $H(Y)$, and X at rate $H(X|Y)$, where $H(\cdot)$ denotes entropy. The Slepian-Wolf theorem ensures that lossless reconstruction can be asymptotically achieved at these rates; however, although the encoder of X does not need to know Y, it does need to know its encoding rate $H(X|Y)$. To address this problem, in [8] it is proposed to employ an embedded coding scheme coupled with a feedback channel, in such a way that the decoder will request as many bits as necessary to decode the data without errors. However, while it is relatively easy to perform embedded DSC, the requirement of a feedback channel is a rather limiting assumption. Alternatively, one should employ some statistical model that allows to estimate the conditional entropy $H(X|Y)$; some preliminary work in this direction has been proposed in [7]. Eventually, when these solutions are not applicable, one can violate the DSC assumption and let the sources communicate to as much an extent as necessary to estimate $H(X|Y)$ (as e.g. in [5]); however, the communication and computation required to estimate $H(X|Y)$ may outweigh the benefits of DSC.

Moreover, typical DSC schemes employ binary coders (e.g., applied to the bit-planes of a transformed signal or image), thus neglecting the correlation among different bit-planes of the signal. We have found that, for lossless coding, neglecting this correlation can lead to a significant performance loss (about 1 bpp for hyperspectral data). The inter-bit-plane correlation is seldom exploited in the literature, and even when it is (see e.g. [9]), it is done for synthetic (e.g., Gaussian) sources, but there is no available realistic probability model to be used in real-world data sources.

In this paper we expand on our previous lossless compression scheme for hyperspectral data, named DSC-CALIC [6], and provide the following contributions. Firstly, we propose a new full-featured statistical model in order to estimate the rates and probabilities necessary to perform DSC encoding; this eliminates the need for any kind of inter-source communication and computing, and is a step towards the full exploitation of DSC potential. Secondly, we take into account the inter-bit-plane correlations by using a multilevel (as opposed to binary) probability model at the DSC decoder; this is exactly the same model used at the encoder, and allows to obtain estimates of the binary probabilities required to initialize and run the binary channel decoders. Thirdly, we upgrade our Slepian-Wolf scheme to a Wyner-Ziv near-lossless compression scheme, i.e. one that minimizes the absolute maximum error, as opposed to the mean squared-error, between the decoded and the original image. This is very important for scientific imagery, where the improved compression ratio comes at the expense of a very small quality degradation, allowing full exploitation of the image information content. Moreover, each band of the compressed file can be decoded progressively.

2. THE PROPOSED NEAR-LOSSLESS DSC-CALIC CODER

We briefly review the operation of DSC-CALIC and describe the modifications for near-lossless and progressive operation; the reader is referred to [6] for additional details.

2.1 Review of CALIC

In the following we briefly review the near-lossless CALIC compression scheme, as it is used to generate the prediction errors encoded by the proposed technique. CALIC is based on non-linear prediction followed by context-based entropy coding. We briefly outline the main steps leading to the generation of the prediction error signal and its entropy coding. We assume that the image to be coded is a band of an hyperspectral data set.

CALIC performs prediction of a pixel value based on a causal neighborhood; the neighborhood is also used to compute contexts for arithmetic coding. It has two operating modes, namely a binary and a continuous mode. The binary mode is only used for those neighborhoods where no more than two distinct intensity values appear; in this case, a special entropy coding mode is triggered. However, we do not employ the binary mode here because it does not generate any prediction error samples but is a sort of direct run-length coding.

		x_6	x_7
	x_3	x_2	x_4
x_5	x_1	x	

Figure 1: Causal neighborhood for prediction in CALIC.

A gradient-adjusted prediction of a pixel is carried out using seven adjacent pixels in its neighborhood (see Fig. 1). Estimates of the horizontal and vertical gradients of the pixel to be predicted x are computed as $d_h = |x_1 - x_5| + |x_2 - x_3| + |x_2 - x_4|$ and $d_v = |x_1 - x_3| + |x_2 - x_6| + |x_4 - x_7|$. According to the values of d_h and d_v, a nonlinear rule is applied in order to obtain an estimate x' of pixel x based on the sharpness of the horizontal and vertical edges.

CALIC also defines contexts, which are used to decrease the prediction error variance, as well as improve the efficiency of entropy coding. It employs 144 texture contents, combined with 4 error energy contexts, leading to overall 576 contexts. In each context, the average value $\overline{e}$ of the prediction error is computed; this average is then subtracted from the prediction error, so as to obtain the bias-free prediction error $Ex = x - x' - \overline{e}$.

In order to achieve near-lossless compression, the prediction error is quantized using a uniform quantizer with odd step-size, $2L_\infty + 1$, yielding the quantization indexes Q_x. The decoder, after inverse quantization, will reconstruct the original image within an absolute maximum error equal to L_∞. In continuous mode, a binary context-based arithmetic coder is used to encode the quantization indexes Q_x (this will be replaced by a Slepian-Wolf coder in the proposed scheme).

2.2 Overview of the proposed scheme

The basic idea of the proposed scheme is to encode each band separately, and to decode the current band using the previous, already decoded band as side information. What is actually entropy-coded are the quantization indexes Q_x of the current band.

As side information, the decoder uses the indexes Q_y of the previous band.

The encoding of one band is done bit-plane by bit-plane, and bit-planes are generated by reading the indexes Q_x in raster-scan order. Bit-planes are not generated by taking the bits of equal degree of all coefficients, but rather using the significance-based ordering ala SPIHT [12], in which a sample x in bit-plane i has a class $C(x,i)$ belonging to one out of three types, namely significance, refinement, and sign.

Each bit-plane is then coded using DSC; we employ syndrome-based coding using low-density parity-check (LDPC) codes of suitable rate. An arithmetic coder can also be used in place of the syndrome coder.

The decoder carries out syndrome decoding (or arithmetic decoding) to recover the bit-planes. It then employs the bit-planes to reconstruct the quantized prediction error samples, and performs inverse quantization and inverse CALIC prediction.

A statistical correlation model is employed to model the signal in order to estimate a few probabilistic parameters that are needed by the encoder and the decoder. The encoder requires to know the conditional entropy of each bit-plane given the corresponding bit-plane in the side information. The decoder needs to know the log-likelihood ratios of the bits to be decoded.

In the following the building blocks of the proposed scheme are described in more detail.

2.3 Statistical model for probability estimation

A good correlation model is the key ingredient in order to easily estimate the bit-plane conditional entropies on the encoder side, as well as to initialize the probabilities of the message passing syndrome decoder. In this paper we propose to employ a Laplacian correlation model; in particular, the following conditional probability density function (pdf) is used:

$$f_{X|Y}(x|y) = \frac{\lambda}{2} e^{-\lambda|x-y|} \quad (1)$$

where the random variables X and Y are, from now on, used to denote the quantized prediction error samples Q_x and Q_y in order to simplify the notation. As a consequence, the variance of X conditioned on Y can be expressed as $\sigma^2_{X|Y} = 2/\lambda^2$, which is a parameter to be estimated. In order to not violate the DSC assumptions, the samples x in the current band cannot be used in the estimation of $\sigma_{X|Y}$; therefore, in our scheme the encoder and decoder obtain an estimate of $\sigma_{X|Y}$ using the two previous bands Y and Z which precede X. In particular, the variance of the difference between the co-located samples of bands Y and Z is used as an estimate of $\sigma^2_{X|Y}$ and the corresponding λ is evaluated and used in Eq. (1). In this way the estimation is causal, and the decoder can perform the same computations as the encoder. As for the mean value y of the conditional distribution, the true value of the sample y is used.

The Laplacian model, which is a continuous-amplitude model, is then used to evaluate the probability of a given sample being equal to 1 in the considered bit-plane n. According to the bit-plane coding approach there are 3 bit classes: significance, sign and refinement bits. The probability that the n-th bit of sample x turns out to be significant is

$$P_S(x,n) = \int_{2^n}^{2^{n+1}} |x| f_{X|Y}(x|y)dx$$

The probability of sample x, which is significant in bit-plane n, being positive can be evaluated as:

$$P_+(x,n) = \int_{2^n}^{2^{n+1}} x f_{X|Y}(x|y)dx$$

Finally, given a partially coded sample x and its current reconstruction $\hat{x}_{n+1}$ in bit-plane $n+1$, the probability of having a refinement bit of value 1 in bit-plane n is

$$P_R(x,n) = \begin{cases} \int_{\hat{x}_{n+1}}^{\hat{x}_{n+1}+2^n} x f_{X|Y}(x|y)dx \ , \ \hat{x}_{n+1} \geq 0 \\ \int_{\hat{x}_{n+1}-2^n}^{\hat{x}_{n+1}} x f_{X|Y}(x|y)dx \ , \ \hat{x}_{n+1} < 0 \end{cases}$$

Note that these integrals have simple closed-form expressions, which are not reported here for brevity. These quantities are used by the encoder to estimate the conditional entropy of bitplane i, $H_i = \sum_x H\left(P_{C(x,i)}(x,i)\right)$, where $C(x,i)$ identifies the bit class of sample x in bit-plane i and $H(p) = -p\log_2(p) - (1-p)\log_2(1-p)$. The bit classes are obtained with bit-plane coding approach ala SPIHT, where all sample are initially tested for significance, when a sample is significant its sign is transmitted, then refinement bits are coded.

Analogously, the bit probabilities computed above are used on the decoder side to initialize the log-likelihood ratios (LLR) of all sample in a certain bit-plane. The value of a given LLR depends on the considered bit-plane and bit class. More details on this aspect are given in Sect. 2.5.

2.4 Encoder

The near-lossless CALIC encoder [10] is employed to generate the quantized prediction error samples X. Slepian-Wolf coding is applied to X as entropy coding stage. In particular, the encoder decomposes X into its bit-planes X_i using the significance/refinement/sign procedure. The equivalent bit-planes Y_i of the prediction error array Y of the previous band will be used by the decoder as side information. For each bit-plane, we define the conditional entropy $H_i = H(X_i|Y_i)$, as well as its estimated value H_i^*, which is computed as explained in Sect. 2.3.

For each bit-plane, depending on the value of H_i^* the encoder selects one out of two possible coding modes, i.e. i) transmit X_i using an arithmetic coder, and ii) employ the DSC mode using syndrome-based coding. If $H_i^* > 0.95$, then the bit-plane is transmitted using an arithmetic coder, since there is no channel code of suitable rate in the database. Otherwise, the DSC mode is triggered and the bit-plane is encoded using syndrome-based coding, i.e. the syndrome Z_i of the array X_i is computed using an LDPC code with number of checks $M > H_i^*$, and is transmitted to the decoder. In particular, we use a database of about 100 irregular LDPC codes [14, 15] with belief-propagation decoding, with code rate between 0.1 and 0.95. For more information on syndrome-based coding the reader is referred to [1]. It is worth noticing that, according to the proposed approach, each bit-plane is coded by a single LDPC; on the contrary, each bit-plane collects bits of different classes and with different probability of being 1 or 0. From the channel coding point of view, a single linear code is being used to face a time-varying channel.

2.5 Decoder

The bit-planes Y_i of the quantized prediction error of the previously decoded band are used as side information to estimate the bit-planes X_i of the quantized prediction error of the current band, which will be combined to form the reconstructed quantization indexes X. Then inverse quantization and CALIC inverse prediction are used to obtain the final decoded band.

The estimation process depends on which coding mode has been employed by the encoder for each bit-plane. If the arithmetic coding mode has been used, an arithmetic decoder is used to extract the bit-plane X_i. If the DSC mode has been used, the decoder runs the iterative message-passing LDPC decoding algorithm to recover X_i. The LDPC decoder takes as inputs the LLRs of the received message, the received syndromes, and the side information Y_i. Unlike classical LDPC decoding, as in [11] we carry out syndrome decoding, attempting to make the decoder converge to an estimated message having exactly the received syndrome Z_i. The most important aspect of LDPC decoding is the LLR estimation. LLRs for a binary asymmetric channel are defined as follows [6]:

$$LLR_y = \log_2 \frac{P(x=1|y)}{P(x=0|y)} = \begin{cases} \log_2 \frac{p_i}{1-p_i} & \text{if } y=0 \\ \log_2 \frac{1-q_i}{q_i} & \text{if } y=1 \end{cases}$$

where $p_i = P(x_i = 1|y_i = 0)$ and $q_i = P(x_i = 0|y_i = 1)$ are assumed to be known. This model neglects the correlation between neighboring bit-planes. To exploit this additional correlation, we employ the statistical model described in Sect. 2.3 to estimate the probabilities of each bit-plane given the side information *and* all the previously decoded bit-planes. Given a bit-plane i, the LLR for each sample x depends on the corresponding bit class $C(x,i)$ and can be expressed as

$$LLR(x,i) = \log_2 \frac{P_{C(x,i)}(x,i)}{1-P_{C(x,i)}(x,i)}$$

where the correlation with the previously decoded bit-planes is carried by the conditional probability $P_{C(x,i)}(x,i)$.

3. EXPERIMENTAL RESULTS

The proposed scheme has been applied to AVIRIS [13] hyperspectral images. The test images have 256 lines, 614 pixels per line, and 224 bands. This leads to a block size of

$N \simeq 10^5$, for which LDPC codes are known to exhibit good error-correction performance. Rather than providing average bit-rates, we show the individual bit-rates for a few bands, which are decoded using the previous band as side information. The bands have been chosen to be roughly uniformly spaced in the instrument wavelength range.

The results obtained on the *Cuprite* scene are reported in Tab. 1 for different values of absolute maximum error L_∞. The table reports the bit-rates, in bits per sample, for various techniques, which have been chosen in order to evaluate how close the proposed technique is to the theoretical performance bounds. In particular, b-AC refers to the proposed algorithm when the DSC mode is disabled, and an ideal binary arithmetic coder is used to encode all bit-planes; i.e., the bit-rate is nothing but the average rate obtained coding each bit-plane as a separate binary source with a rate equal to its entropy. The technique labeled as b-DSC consists in the proposed technique, in which the LDPC entropy coder is replaced by an ideal *binary* Slepian-Wolf coder; in other terms, the bit-rate is the average rate obtained coding each bit-plane as a separate binary source with a rate equal to its conditional entropy given the side information. This technique serves as reference to assess the best result that could be achieved by a coder that neglects the inter-bit-plane correlations. The column labeled as H_C reports the conditional entropy of the multilevel (as opposed to binary) source, and is used here as performance bound for the proposed multilevel LDPC Slepian-Wolf coder. Finally, the technique labeled as LDPC is the proposed technique with a true LDPC coder and decoder.

Several remarks can be made analyzing the results in Tab. 1. It can be seen that the performance difference between b-AC and b-DSC is significant (about 1 bpp); this highlights the potential of DSC-based techniques to capture the correlation between adjacent bands. By comparing b-DSC and H_C, it can be noted that the correlation between adjacent bit-planes is very significant, since the binary algorithm exhibits an average performance loss in excess of 1 bpp. The practical LDPC-based scheme performs significantly better than b-DSC, but is still far from the performance bound of H_C. This is due to the fact that we are using a binary channel coder outside the capacity region of the binary channel (independent transmission and decoding of bit-planes), attempting to reach the (higher) capacity of the multilevel channel by providing a more accurate statistical model. Since this channel coder has not been developed nor optimized for multilevel transmission, a performance gap has to be expected when it is used in the multilevel context. From Tab. 1, this gap turns out to be between 0.7 and 1 dB, whereas on the same data we have found that the binary LDPC code is as close as 0.05 bpp to the capacity of the binary channel. However, there is still a significant improvement (up to about 0.5 bpp) with respect to the ideal performance bound for the binary channel. Similar results, reported in Tab. 2, have been obtained on the *Jasper Ridge* scene.

For comparison, in Tab. 3 we also report the bit-rates achieved by CALIC [10] on the same scenes and in the same conditions as above. As can be seen, the performance of CALIC is slightly better than that of b-AC thanks to the context modeling. However, both the practical proposed technique and its performance bound exhibit significantly lower bit-rates, witnessing that DSC is indeed able to provide significantly better performance than lossless 2D techniques, although there is still some work to be done in the design of multilevel Slepian-Wolf coders, in order to get more close to the performance bounds.

Table 1: Bit-rates obtained on the *Cuprite* scene.

L_∞	Band	b-AC	b-DSC	H_C	LDPC
1	20	7.52	6.39	5.17	5.98
1	60	5.66	4.90	3.88	4.75
1	95	5.88	4.92	3.82	4.55
1	130	5.37	4.37	3.26	3.98
1	180	4.80	4.33	3.46	4.19
3	20	6.32	5.20	4.06	4.69
3	60	4.50	3.80	2.76	3.76
3	95	4.72	3.82	2.73	3.69
3	130	4.21	3.25	2.22	3.09
3	180	3.64	3.21	2.37	3.27

Table 2: Bit-rates obtained on the *Jasper Ridge* scene.

L_∞	Band	b-AC	b-DSC	H_C	LDPC
1	20	7.63	6.21	4.97	5.93
1	60	7.67	6.22	4.88	5.82
1	95	7.16	5.56	4.26	5.06
1	130	6.36	4.55	3.26	3.99
1	180	5.56	4.59	3.53	4.34
3	20	6.44	5.04	3.90	4.79
3	60	6.47	5.05	3.78	4.56
3	95	5.97	4.42	3.16	3.95
3	130	5.18	3.49	2.27	3.16
3	180	4.39	3.55	2.47	3.47

Table 3: Bit-rates obtained by CALIC on the *Cuprite* and *Jasper Ridge* scenes.

L_∞	Band	Cuprite	Jasper
1	20	7.40	7.32
1	60	5.44	7.56
1	95	5.83	7.02
1	130	5.11	6.09
1	180	4.24	5.14
3	20	6.20	6.09
3	60	4.24	6.33
3	95	4.62	5.80
3	130	3.91	4.88
3	180	3.06	3.60

Finally, in order to further validate the proposed statistical model, we have estimated the multilevel conditional entropy by using our conditional pdf model (see Eq. 1) inside the definition of conditional entropy:

$$H(X|Y) = -\sum_x \sum_y f_{X|Y}(x|y) f_Y(y) \log f_{X|Y}(x|y).$$

It turns out that our estimate is typically within 0.05 and 0.1 bpp from H_C when the true and estimated values of $\sigma_{X|Y}$ are employed respectively, thus witnessing the accuracy of the proposed model.

4. CONCLUSIONS

In this paper we have proposed a compression technique for hyperspectral data based on DSC. Notable features are near-lossless reconstruction, and the use of a statistical model for encoder and decoder probability estimation.

We have found that the statistical model provides accurate estimates of the required probabilities, eliminating the need of inter-band communication (or a feedback channel).

The same model can be used to improve the LLR estimation for the LDPC decoder, in an attempt to boost its performance and achieve a rate equal to the conditional entropy of the multilevel source. We have found that the proposed approach does yield a nice performance improvement. However, since the LDPC coder is a binary channel coder, it still exhibits a performance gap with respect to an ideal multilevel Slepian-Wolf coder.

REFERENCES

[1] Z. Xiong, A.D. Liveris, and S. Cheng, “Distributed source coding for sensor networks,” *IEEE Signal Processing Magazine*, vol. 21, no. 5, pp. 80–94, Sept. 2004.

[2] D. Slepian, J.K. Wolf, “Noiseless coding of correleted information sources,” *IEEE Transactions on Information Theory,* vol 19, pp. 471-480, July 1973.

[3] A. Wyner and J. Ziv, “The rate-distortion function for source coding with side information at the decoder,” *IEEE Transactions on Information Theory*, vol. 22, pp. 1–10, Jan. 1976.

[4] R. Zamir, “The rate loss in the Wyner-Ziv problem,” *IEEE Transactions on Information Theory*, vol. 42, no. 6, pp. 2073–2084, Nov. 1996.

[5] R. Puri and K. Ramchandran, “PRISM: a “reversed” multimedia coding paradigm,” in *Proc. of IEEE International Conference on Image Processing*, 2003.

[6] A. Nonnis, M. Grangetto, E. Magli, G. Olmo, and M. Barni, “Improved low-complexity intraband lossless compression of hyperspectral images by means of slepian-wolf coding,” in *Proceedings of IEEE International Conference on Image Processing*, 2005.

[7] N.M. Cheung, H. Wang, and A. Ortega, “Correlation estimation for distributed source coding under information exchange constraints,” in *Proceedings of IEEE International Conference on Image Processing*, 2005.

[8] B. Girod, A. Aaron, S. Rane, and D. Rebollo-Monedero, “Distributed video coding,” *Proceedings of the IEEE*, vol. Special Issue on Advances in Video Coding and Delivery, no. 1, pp. 71–83, Jan. 2005.

[9] S. Cheng, Z. Xiong, “Successive refinement for the Wyner-Ziv problem and layered code design”, *IEEE Transactions on Signal Processing*, vol. 53, n. 8, pp. 3269-3281, Aug. 2005.

[10] X. Wu and N. Memon, “Context-based, adaptive, lossless image coding,” *IEEE Transactions on Communications*, vol. 45, no. 4, pp. 437–444, Apr. 1997.

[11] A. Liveris, Z. Xiong and C. Georghiades, “Compression of binary sources with side information at the decoder using LDPC codes,” *IEEE Communications Letters*, vol. 6, pp. 440-442, October 2002.

[12] A. Said and W. A. Pearlman, “A new, fast, and efficient image codec based on set partitioning in hierarchical trees”, *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, no. 3, June 1996, pp. 243-250.

[13] M.K. Simon, D. Divsalar, “Lossless Compression of AVIRIS images,” *IEEE Transactions on Image Processing*, vol. 5, pp. 713-719, May 1996.

[14] R. Gallager, “Low Density Parity Check Codes”, *MIT Press*, 1963.

[15] D. MacKay, “Good error-correcting codes based on very sparse matrices,” *IEEE Transactions on Information Theory*, vol. 45, pp. 399-431, March 1999.