Analysis of Trellis Quantization for Near-Lossless Image Coding[1]

Hannes Hartenstein
Institut für Informatik
Universität Freiburg
Am Flughafen 17
79110 Freiburg, Germany
hartenst@informatik.uni-freiburg.de

Xiaolin Wu
Dept. of Computer Science
Middlesex College 368
University of Western Ontario
London, Ontario, Canada N6A 5B7
wu@csd.uwo.ca

Abstract

In this paper we discuss near-lossless image compression by trellis quantization (TQ) coupled with predictive coding. Here near-lossless coding means L_∞-constrained coding, i.e., each pixel can be altered by at most τ grey values, where τ is a pregiven threshold. In a predictive coding scheme the simplest way to guarantee the L_∞ bound is to perform a uniform quantization of the prediction residues with quantizer bin size matching the required tolerance. The near-lossless version of CALIC [1] uses this mechanism. But since quantization is inside the DPCM loop changing a pixel value at the current position affects subsequent predictions of forthcoming pixels. The trellis quantization scheme proposed by Ke and Marcellin [2] tries to take into account those global implications of quantization. However, despite its high computational cost the TQ scheme is inferior to the near-lossless CALIC. In this paper we reexamine the trellis quantization approach for near-lossless image coding, and point out its weaknesses that render TQ ineffective. We also incorporate into TQ context modeling techniques inspired from CALIC. We report coding results for various standard test images.

1 Introduction

Many important applications such as medical imagery, satellite imagery, and numerical weather simulations would opt for lossless image compression. But the compression ratios achievable by lossless coding are usually not sufficient. Thus, errors have to be tolerated in order to obtain higher compression ratios. Since the image data in the applications mentioned above is processed by computers and/or highly trained experts, distortion measures based on visual impression are inappropriate. In these situations each pixel value is of equal importance. Therefore, the distortion in each pixel has to be controlled. This naturally leads to L_∞-constrained compression, where each pixel is only altered by at most τ grey levels, where τ is a user-specified threshold. For example, in the context of the weather simulation data, the value of τ depends on the accuracy of measurements of the corresponding physical modality. For small τ L_∞-constrained image coding is commonly referred to as near-lossless compression. Of course, the value τ can be set in a region-dependent way. In this paper we consider only one global τ parameter.

[1] An abstract of this report is published in: *Proceedings of the IEEE Data Compression Conference 1998*, J. Storer, M. Cohn (eds.), Computer Society Press, Los Alamitos, 1998

A predictive lossless image coder can easily be modified to serve as a near-lossless coder. Instead of the lossless coding of the prediction residues, the residues are uniformly quantized with the required precision. For example, the near-lossless version of CALIC [1] works in that fashion.

Trellis quantization is a technique to minimize the bit rate under a given L_∞ bound. Namely, we replace uniform memoryless scalar quantization of prediction residues by trellis quantization as proposed by [2]. Instead of a simple nearest neighbor quantizer mapping on a pixel-by-pixel basis, trellis quantization tries to factor in the global impact of individual quantization errors on predictions and on bit rate. In [2] Ke and Marcellin used dynamic programming technique to design a trellis quantizer aiming to minimize conditional entropy in a few simple coding contexts. Moayeri has presented a combination of trellis quantization and a generalized run length encoding scheme in [4]. Unfortunately, trellis quantization has very disappointing performances in near-lossless image coding. It did not offer, in spite of its high computational complexity, appreciable compression gains over uniform quantization in the framework of predictive coding in our experiments [5]. In fact, results obtained by TQ have been much worse compared to near-lossless CALIC. For tolerances $\pm 1, \pm 2$ there was not even a gain over standard DPCM combined with uniform quantization. In this paper we report on our work to improve the TQ scheme of [2] and analyze the major drawbacks of TQ. In section 2 we review some theoretical results related to TQ-based near-lossless image coding from our paper [6]. Several variations to the original scheme are presented in section 3, where we also give coding results for various standard test images from the ISO/IEC test set. Section 4 concludes the paper with an evaluation of the trellis quantization approach and with some open problems.

2 Theoretical Results

In this section we state the trellis quantization problem of [2] in the context of one dimensional signals. We answer the question of how large the maximal coding gain can be. The code for a finite one-dimensional signal $s \in \mathbb{Z}^n$ has to represent a reconstruction signal $\hat{s} \in \mathcal{N}_\tau(s)$, where

$$\mathcal{N}_\tau(s) := \{t \in \mathbb{Z}^n | \; \|t - s\|_\infty \leq \tau\},$$

with $\|t - s\|_\infty = \max_{i \in \{1..n\}} |t_i - s_i|$. $\mathcal{N}_\tau(s)$ can be seen as the set of all left-to-right paths in a trellis as depicted in Figure 1. Let D be the difference operator $D : \mathbb{Z}^n \to \mathbb{Z}^{n-1}$, $s = (s_1, \ldots, s_n) \mapsto (s_2 - s_1, \ldots, s_n - s_{n-1})$. The distribution $P(s)$ of a signal $s \in \mathbb{Z}^n$ is given by $p_j(s) = \frac{1}{n} |\{i | s_i = j\}|, j \in \mathbb{Z}$, and we define the entropy of s by $H(s) := H(P(s)) = -\sum_{j \in \mathbb{Z}} p_j(s) \log_2(p_j(s))$.

We are interested in finding the reconstruction signal in the neighborhood $\mathcal{N}_\tau(s)$ that can be coded most efficiently using differential coding, i.e., we are faced with the following optimization problem: find $t^* \in \mathcal{N}_\tau(s)$ such that $H(D(t^*)) = \min_{t \in \mathcal{N}_\tau(s)} H(D(t))$. In contrast to the standard bit allocation problem (see e.g. [7]) here the distortion is fixed and the number of bits has to be minimized. The modification of the signal value s_i not only changes the residue $s_i - s_{i-1}$ but also affects the prediction for the residue $s_{i+1} - s_i$, and hence has effects on the coding of the rest of the signal. This chain effect makes the above

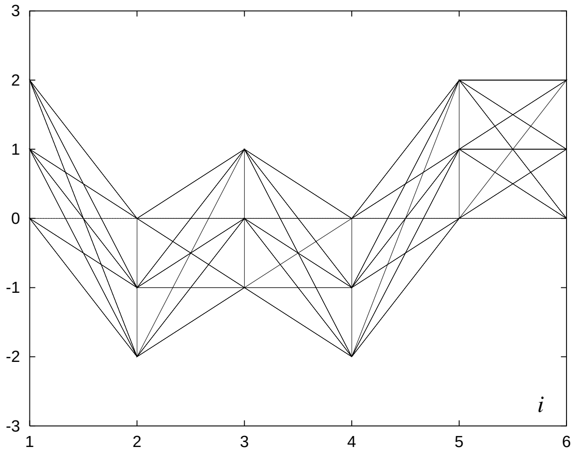

Figure 1: Each left-to-right path corresponds to an element of $\mathcal{N}_\tau(s)$ with $s = (1, -1, 0, -1, 1, 1)$ and $\tau = 1$.

optimization problem very hard from the computational complexity point of view.

Since one cannot find the optimal solution in a practical amount of time, it is useful to have a lower bound for $\min_{t\in\mathcal{N}_\tau(s)} H(D(t))$. When there is no prediction involved, i.e., if one is interested in $\min_{t\in\mathcal{N}_\tau(s)} H(t)$, the following inequality holds:

$$H(t) \geq H(s) - \log_2(2\tau + 1), \tag{1}$$

for any signal $s \in \mathbb{Z}^n$ and $t \in \mathcal{N}_\tau(s)$. This is a consequence of a theorem presented by Gormish and Gill in [8].

This result can be extended to the case of differential coding:

Theorem 1 *For all $t \in \mathcal{N}_\tau(s)$ the following inequality holds:*

$$H(D(t)) \geq H(D(s)) - \log_2(4\tau + 1). \tag{2}$$

Details of the proofs for inequalities (1) and (2) are given in [6]. The bound of the theorem can be attained, e.g., with the signal in Figure 1 and $\tau = 1$ (if one takes the horizontal line reconstruction). Therefore, the lower bound is indeed tight. Note that by simply using a uniform quantizer with bin size 3 for the prediction residues one obtains the following sequence of quantized residues: -3, 3, -3, 3, 0. This sequence has an entropy value of 1.38, whereas the optimal sequence (the horizontal line reconstruction) has entropy value 0. The bound of (2) is derived from analyzing the distribution of the differences between neighboring samples of the original signal. However, while this result represents the best possible bound, in most cases it cannot be attained since the actual order in which the signal values are given is important in practice but not used in the derivation of the theorem.

It is easy to see that there are also signals such that the bound of inequality (2) can be attained by uniform quantization of the prediction residues. This is also true even when the signal values are first quantized and then the differences are losslessly coded, as one can see from the signal in Figure 1.

The example above and the lower bound indicate that it could be worth trying to find good solutions to the optimization problem. In the next section we review the iterative optimization procedure proposed by Ke and Marcellin in [2].

3 Original Approach and Variations

Let us first explain the basics of the trellis quantization scheme for the coding of 1-dimensional signals:

1. Input: signal $s \in \mathbb{Z}^n$, tolerance $\tau > 0$.
2. Initialize the probability distribution $P^* = (p_j^*)_j$ for the residue symbols.
3. Find optimal reconstruction for the current probability distribution P^*, i.e., find $\hat{s} \in \mathcal{N}_\tau(s)$ such that $-\sum p_j \log p_j^*$ is minimized, where $(p_j)_j = P(D(\hat{s}))$.
4. Update P^* with the distribution $P(D(\hat{s}))$. Stop when the entropy gain is smaller than a pregiven threshold, otherwise go to 3.

In this procedure the entropy is monotonically decreasing to a local optimum. Step 3 corresponds to a shortest path problem and can be solved in linear time (depending on the signal length) using a dynamic programming approach.

Let us now switch to the coding of images using a predictive scheme where a pixel $x_{i,j}$ is predicted via the planar predictor, i.e., $x_{i,j}$ is predicted by $\hat{x}_{i,j-1} + \hat{x}_{i-1,j} - \hat{x}_{i-1,j-1}$. Here, $\hat{x}$ denotes a pixel of the reconstruction image. Because of the fact that the predictor extends itself over two image rows, the coding of one row affects the coding of the next row. Therefore, step 3 in the above algorithm gets more complicated, in fact it cannot be computed optimally in a practical amount of time (we conjecture that this problem is NP-hard). Instead, a *greedy* scheme is employed where the rows are successively and independently quantized using the 1-dimensional algorithm explained above.

For simplicity reasons we have only used one probability distribution P^* so far. The original scheme uses conditional probabilities based on contexts to improve the coding performance. The contexts are based on the magnitude of the local gradient $(\hat{x}_{i,j-1} - \hat{x}_{i-1,j-1}) + (\hat{x}_{i-1,j} - \hat{x}_{i-1,j-1})$. Figure 2 gives an example of the evolution of a distribution for a fixed coding context. Compression results are given in Table 1 for the 512×512 8-bit standard test images Lenna and Barbara, as well as for the 8-bit images Aerial2 (2048×2048), Finger (fingerprint image 512×512), and Us (ultrasound image 448×512) from the ISO test set. The number of contexts in the TQ algorithm is set to 9. For reasons of comparison we also give coding results for CALIC and for a DPCM scheme that uses the same predictor as the trellis quantizer but only performs a uniform quantization on the prediction residues. Most of the times the trellis quantization approach is even slightly worse than the DPCM scheme, and its coding performance is much worse compared to the performance of CALIC.

A lossless coding of the image Lenna using a DPCM scheme with predictor and contexts as in the trellis quantization algorithm results in 4.54 bpp coding costs, about 0.5 bpp more than the state-of-the-art lossless coders need. For tolerances $\pm 1, ..., \pm 5$ this gap in

	$\tau = 1$			$\tau = 2$			$\tau = 3$			$\tau = 5$		
	TQ	DP	CA	TQ	DP	CA	TQ	DP	CA	TQ	DP	CA
lenna	3.09	3.00	2.56	2.47	2.39	1.92	1.93	1.97	1.53	1.37	1.49	1.07
	50.0	49.9	49.9	45.3	45.1	45.1	42.5	42.1	42.3	38.8	38.2	38.8
barb	4.10	3.92	3.07	3.40	3.30	2.63	2.91	2.83	2.02	2.23	2.26	1.55
	50.1	49.9	49.9	45.3	45.1	45.2	42.3	42.1	42.3	38.7	38.1	38.6
aerial2	4.13	3.94	3.63	3.31	3.25	2.93	2.94	2.78	2.48	2.29	2.18	1.92
	49.9	49.9	50.0	45.3	44.9	45.2	42.4	41.8	42.2	38.3	38.3	38.2
finger	4.22	4.05	3.82	3.46	3.34	3.11	3.01	2.88	2.66	2.38	2.31	2.11
	50.1	49.8	49.9	45.2	45.1	45.2	42.3	42.1	42.2	38.3	38.1	38.2
us	1.90	1.87	1.53	1.58	1.58	1.22	1.37	1.30	1.04	1.09	1.18	0.82
	50.6	51.0	51.8	46.0	45.5	47.7	42.9	44.4	44.6	39.1	36.1	40.1

Table 1: Comparison of coding results for trellis quantization (TQ), DPCM plus uniform quantization (DP), and CALIC. Given are bits-per-pixel and PSNR values.

bit rates gets only marginally smaller. This indicates that the predictor and contexts should be improved and that more adaptivity is needed. In the next subsections these issues are discussed.

3.1 Two-row trellis quantization

In [2] the trellis quantizer was designed for only the current row of pixels with the previous row fixed. In reality, however, the prediction robustness heavily depends on adjacent pixels in the previous row, and the propagation of quantization distortions is two-dimensional. With these considerations we increase the scope of the trellis from one row to two rows, and minimize the total code length of two rows of pixels under the L_∞ constraint via dynamic programming and a similar iterative process as in [2].

A one-row trellis has n stages and $2\tau + 1$ states at each stage (we neglect the boundary problems). Thus, the complexity of the one row optimization is proportional to $n \cdot (2\tau+1)^2$. A two-row trellis has $(2\tau+1)^2$ states at each stage (see Figure 3). In general, a k-row trellis-based joint optimization would require $n \cdot (2\tau + 1)^{2k}$ operations. We emphasize that even the two-row joint optimization has a very high computational complexity. We implemented the two-row trellis primarily for benchmarking purposes. Results for tolerances $\pm 1, \pm 2, \pm 3$ are given in Table 2. To our disappointment the two-row trellis only leads to a marginal coding gain. We also tested *overlapping* optimization, i.e., a two-row joint optimization is performed but only the first row is fixed and the next two-row optimization is performed with the second and third row, and so on. This leads to similar results.

3.2 Progressive probability update

With the iterative TQ design scheme of [2] one obtains probability distributions that are adaptively determined for a particular input image. These distributions have to be sent as side information in order to have the corresponding entropy coders available at the decoder

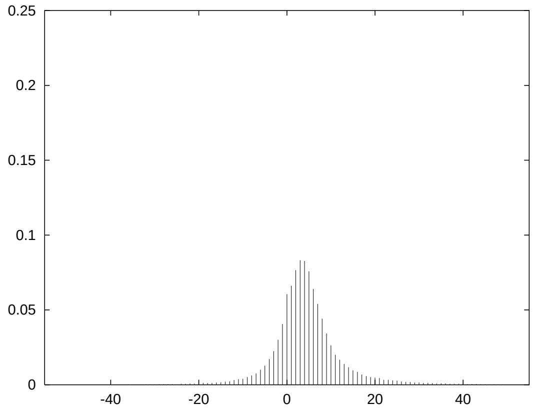

1 iteration

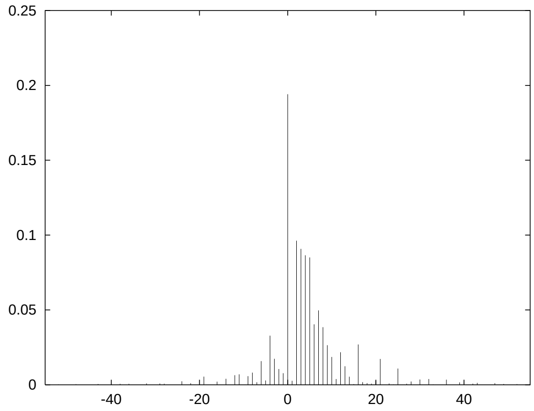

3 iterations

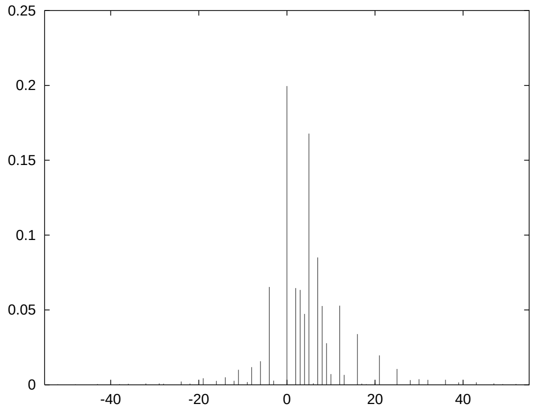

5 iterations

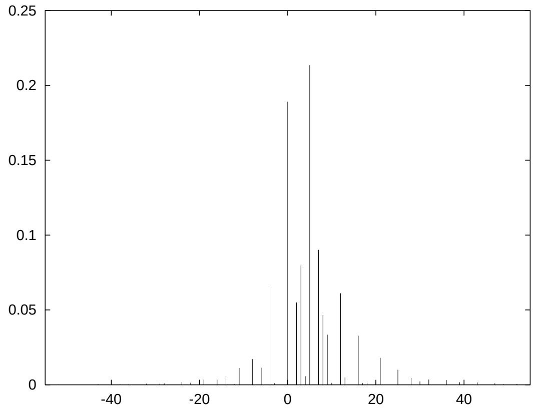

7 iterations

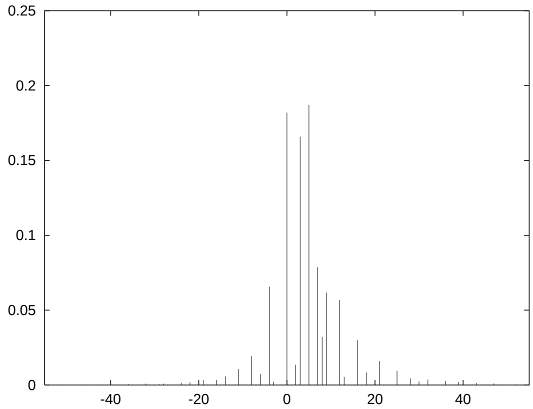

11 iterations

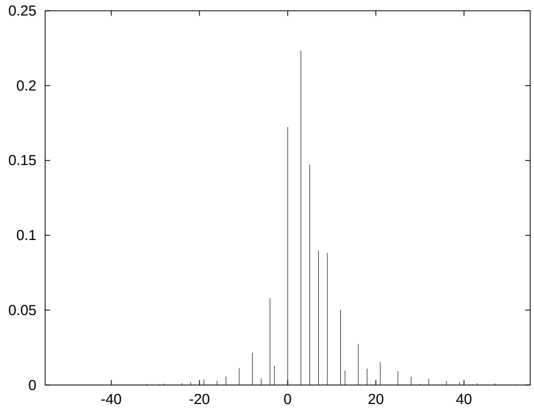

19 iterations

Figure 2: Shown is the evolution of a (conditional) distribution toward convergence. This particular example is for tolerance ± 1.

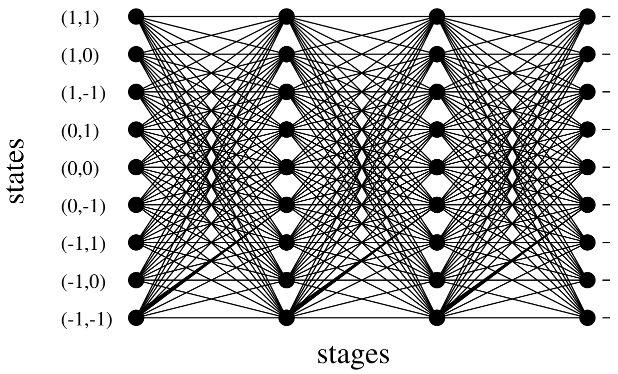

Figure 3: Schematic representation of the trellis for the two-row optimization (for tolerance ± 1). At stage j the state labeled (m, n) corresponds to the reconstruction values $(\hat{x}_{i,j} + m, \hat{x}_{i+1,j} + n)$. The edge weights are given by the coding costs of the corresponding prediction residuals.

	$\tau = 1$		$\tau = 2$		$\tau = 3$	
	1-row-opt.	2-row-opt	1-row-opt.	2-row-opt	1-row-opt.	2-row-opt
lenna	3.09	3.08	2.47	2.37	1.93	1.85
finger	4.22	4.19	3.46	3.44	3.01	3.00

Table 2: Comparison of 1-row optimization versus 2-row optimization. All numbers are bits per pixel values.

side. The cost for the side information is relatively low, e.g. lower than 0.05 bpp when 9 coding contexts are used. Another possibility is to use the iterative procedure on a training set of images to determine the distributions of quantized residues for entropy coding. The resulting distributions are then made available to and used by both the encoder and decoder. This eliminates the use of side information. In [2] the second method is preferred over the adaptive scheme because it does not require the expensive iterative design process. In our experiments the fixed scheme has shown a worse performance than the adaptive scheme. This will be the case as long as the input image has different statistics from that of the training set. Furthermore, image statistics are hardly stationary. This is why most of the successful lossless image coders update the distributions after coding each symbol to follow non-stationarity of the source. With the TQ approach one cannot update the distributions after each symbol, since the dynamic programming process requires fixed conditional distributions in minimizing the bit rate with respect to the distributions. Nevertheless, we can update the probability distributions on a row by row basis with trellis quantization of prediction residues. Namely, we can update source statistics with all residue terms in a row once the row has been determined by the trellis quantizer. The effects of row-by-row updating are graphically illustrated in Figure 4.

3.3 Lossy plus near-lossless coding

When one compares CALIC to the TQ approach it is obvious that CALIC uses much more sophisticated context modeling and prediction schemes. Is it possible to incorporate such mechanisms into TQ? The difficulty lies in the high computational complexity. Assume that we like to use larger contexts in the optimization of trellis quantization. For example, for coding the pixel $x_{i,j}$ we include an additional term $x_{i,j-2}$ into the context calculation.

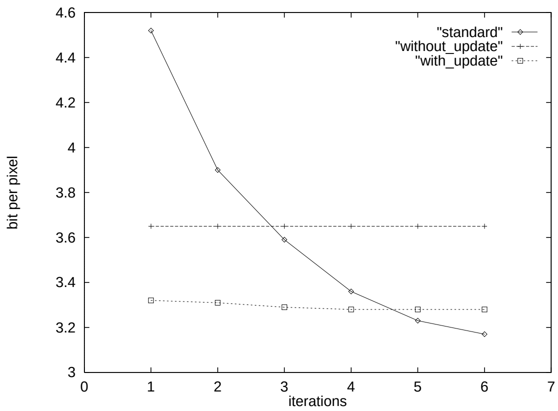

Figure 4: *Standard* refers to the original algorithm where the distributions are updated after each iteration. For the two other curves the algorithm started with a set of trained distributions and proceeded with/without a probability update after each row.

Then instead of $n(2\tau+1)^2$ one has $n(2\tau+1)^3$ operations for the one-row optimization. It is this exponential growth in computation time that prevents the use of more pixels of previous lines.

In our study [5] we have experimented with lossy plus near-lossless residue encoding. A low bit-rate lossy version of the image being encoded is produced using, e.g., the SPIHT algorithm [9]. Then the difference image between original and lossy version is near-losslessly coded, for example with the trellis quantization approach. Now since the lossy version is fixed and available, one can use it to enrich modeling contexts without making the trellis quantization optimization more complex. For example, we have used the magnitude of $|\tilde{x}_{i-1,j}-\tilde{x}_{i,j}|+|\tilde{x}_{i+1,j}-\tilde{x}_{i,j}|+|\tilde{x}_{i,j-1}-\tilde{x}_{i,j}|+|\tilde{x}_{i,j+1}-\tilde{x}_{i,j}|+2|\hat{x}_{i-1,j}|+2|\hat{x}_{i,j-1}|$ to determine the context for position $x_{i,j}$. Here, $\tilde{x}$ is a pixel from the lossy version and $\hat{x}$ is a pixel from the reconstruction image. Coding results of the hybrid scheme are given in table 3. The wavelet transform can capture global features very efficiently in terms of higher order modeling with a relatively small number of parameters. This can explains why the hybrid approach improves the results for the lenna image but has no positive affect for the coding of the fingerprint image.

The use of higher order linear predictors in conjunction with trellis quantization leads to the same problem of computational intractability as the use of higher order contexts discussed above. Of course, the planar predictor can be replaced by a linear predictor that is designed off line and is optimal in the mean square error sense. However, since the optimized predictor is determined independent of the trellis quantization optimization process, it did not lead to any coding gains in our experiments. In fact, the optimized predictor determined for the difference image has only small coefficients. Thus, with high probability the modifications of the pixel values do not change the next prediction value. So we end up with a situation where the coding of the residue and the prediction of the next sample are independent of each other. In this case trellis quantization is of no use and the optimal strategy is to simply apply a uniform quantizer.

3.4 Classified adaptive entropy coding

The results of the previous sections have indicated that the inferior performance of the TQ approach lies probably in the choice of predictor and contexts itself. In this section we present results of a new near-lossless coding method that is based on the lossless coder pre-

	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$	$\tau = 5$
lenna	3.05	2.33	1.79	1.46	1.17
finger	4.25	3.51	3.05	2.72	2.46

Table 3: Results (in bits-per-pixel) for the *lossy plus near-lossless coding* approach.

	$\tau = 1$		$\tau = 2$		$\tau = 3$		$\tau = 5$	
lenna	2.71	49.9	2.08	45.1	1.70	42.2	1.19	38.7
barbara	3.35	49.9	2.69	45.1	2.28	42.2	1.75	38.4
aerial2	3.77	49.9	3.05	45.1	2.62	42.1	2.03	38.2
finger	4.14	49.9	3.41	45.1	2.95	42.1	2.35	38.1
us	1.79	50.3	1.43	46.5	1.21	41.4	0.95	39.7

Table 4: Coding results (bits-per-pixel and PSNR values) for the *classified adaptive entropy coding* scheme.

sented in [10, 11] to show the advantages of a more adaptive prediction/context-modeling scheme. The method works essentially as follows: a set of 16 entropy coders is maintained. The residue image is partitioned into blocks of size 8×8, and for each block the optimal entropy coder from the set of coders is chosen. The main question is how to determine the optimal set of entropy coders. In [10, 11] an iterative optimization method that converges to a local optimum is proposed. In our tests we have uses a classification of the residual blocks based on the variance. For each class of blocks the optimal entropy coder is determined. The improvements obtained with the iterative optimization procedure were rather small. For the prediction we have used the gradient-adjusted predictor of the CALIC method [3]. Coding results are given in Table 4. With the exception of the fingerprint image, all coding results are quite close to the results obtained with CALIC.

4 Conclusions

We have presented some theoretical remarks on the maximal entropy gain that can be expected when using differential coding and given an L_∞ bound. We have discussed the weaknesses of the trellis quantization approach and have tried to improve the performance by introducing various variations to the original scheme of [2]. But even with the introduced enhancements the TQ approach has so far not been able to outperform near-lossless CALIC in general.

We have two explanations for the ineffectiveness of trellis quantization for near-lossless image compression. First, it is computationally intractable to include the pixels to be determined by the TQ process in the modeling contexts for entropy coding of the residues. TQ has to fix all conditional probabilities in the dynamic programming process that determines the quantization of all pixels for minimum rate. But in uniform or any memoryless quantization, it is computationally feasible to use high order contexts to model and code the source. Furthermore, without trellis one can easily update the conditional probabilities on a pixel-by-pixel basis, adapting to changing statistics of the source much faster.

Second, after prediction and particularly after sophisticated context modeling of the residues as in near-lossless CALIC, the residue sequence in a conditioning state of the entropy coder becomes almost i.i.d., leaving very little structures for the trellis quantizer to exploit.

An interesting open question is whether one can prove that for certain sources (stochastic processes), when coded by a predictive scheme, the uniform quantization of the residues is indeed optimal or at least nearly optimal for near-lossless compression.

References

[1] Wu, X., Bao, P., Choi, W., *L_∞-constrained high-fidelity image compression via adaptive context modeling*, Proc. of the IEEE Data Compression Conference, Snowbird, Utah, 1997

[2] L. Ke and M. W. Marcellin, *Near-lossless image compression: minimum-entropy, constrained-error DPCM*, Proc. IEEE International Conference on Image Processing, Washington, D.C., Oct. 1995.

[3] Wu, X., Memon, N., *Context-based, adaptive, lossless image coding*, IEEE Trans. Communications, vol. 45, no. 4, April 1997

[4] Moayeri, N., *A near-lossless trellis-searched predictive image compression system*, Proc. IEEE International Conference on Image Processing, Lausanne, Sep. 1996

[5] Hartenstein, H., Herz, R., Saupe, D., *A comparative study of L_∞-distortion limited image compression algorithms*, Proc. Picture Coding Symposium, Berlin, Sep. 1997.

[6] Hartenstein, H., Saupe, D., *On entropy minimization for near-lossless differential coding*, submitted to IEEE Communications Letters

[7] Trushkin, A. V., *Bit number distribution upon quantization of a multivariate random variable*, Problems in Information Transmission, 16, 1980

[8] M. J. Gormish and J. T. Gill, *Discrete minimum entropy quantization*, IEEE Trans. Image Processing, vol. 4, pp. 1314–1317, Sep. 1995.

[9] Said, A., Pearlman, W. A., *A new fast and efficient implementation of an image codec based on set partitioning in hierarchical trees*, IEEE Transactions on Circuits and Systems for Video Technology, vol. 6, June 1996, pp. 243–250

[10] Golchin, F., Paliwal, K., *Minimum-entropy clustering and its application to lossless image coding*, Proc. IEEE International Conference on Image Processing, Santa Barbara, Oct. 1997.

[11] Golchin, F., Paliwal, K., *Classified adaptive prediction and entropy coding for lossless coding of images*, Proc. IEEE International Conference on Image Processing, Santa Barbara, Oct. 1997.