INTERLEAVED QUANTIZATION FOR NEAR-LOSSLESS IMAGE CODING

Ignacio Ramírez

Universidad de la República (Uruguay). Facultad de Ingeniería

Abstract. Signal level quantization, a fundamental component in dig-ital sampling of continuous signals such as DPCM, or in near-lossless predictive-coding based compression schemes of digital data such as JPEG-LS, often produces visible banding artifacts in regions where the input signals are very smooth. Traditional techniques for dealing with this issue include dithering, where the encoder contaminates the input signal with a noise function (which may be known to the decoder as well) prior to quantization. We propose an alternate way for avoiding banding artifacts, where quantization is applied in an interleaved fashion, leaving a portion of the samples untouched, following a known pseudo-random Beroulli sequence. Our method, which is sufficiently general to be ap-plied to other types of media, is demonstrated on a modified version of JPEG-LS, resulting in a significant reduction in visible artifacts in all cases, while producing a graceful degradation in compression ratio.

1 Introduction

Predictive coding is one of the oldest, yet still most popular tools for signal sampling, coding and compression [1,2]. The basic idea is to encode data sequentially so that the value of a new sample is encoded differentially with respect to a causal prediction computed from previously encoded samples. This helps in decorrelating the signal, and the prediction errors to be encoded usually exhibit a distribution that is sharply peaked at 0 [3], for which efficient entropy coding methods such as Golomb-Rice are available [4,5,6].

The usual method for improving compression rates in predictive coding is to allow a small distortion in the encoded signal by quantizing the prediction errors in steps of size $\varDelta = 2\delta + 1$, where δ is a positive integer [7]. For small values of δ, this method is often referred to as *near-lossless* compression, since the maximum per-sample distortion is guaranteed to be no more than δ.

There is, however, an important drawback of quantization which applies to all forms of digital representation of signals, including PCM, DPCM, and modern predictive coding: when the signals being encoded are very smooth, the quantization error sequence is highly correlated, creating “bands” or “staircases” which significantly affect the perceived quality of the reconstructed signals (see figures 1 and 3). In the case of predictive coding, this has the additional effect of being fed back into the predictor itself, creating more complex and perhaps even more annoying artifacts.

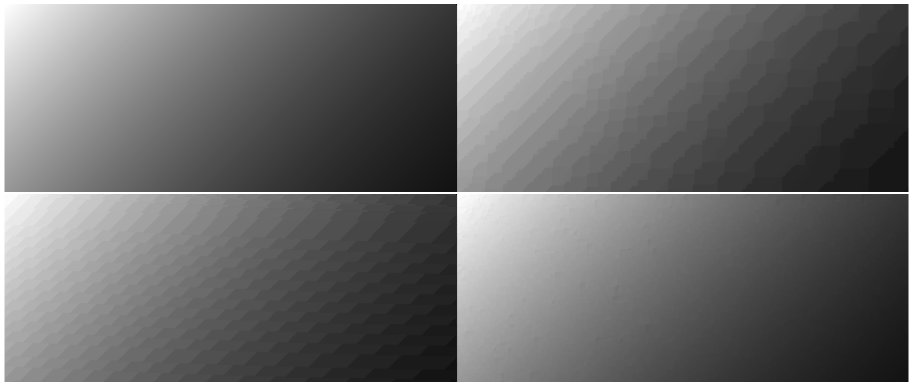

Fig. 1: 2D banding effect. This artifact is more evident (and detrimental to the visual quality) in areas with slowly varying intensity, such as the toy example shown in this figure. Top to bottom, left to right: undistorted image, proposed method (IQ) with $\delta = 4$ and no dithering, JPEG-LS with $\delta = 4$, IQ with $\delta = 4$ and $p = 0.9$ (The differences in the quantized output between JPEG-LS and IQ are due to a modified run length coding method in the latter). *Note: this figure is best appreciated on a computer screen.*

The technique of *dithering* was originally introduced in [8] precisely for reducing banding effects due to quantization in digital signal coding. In short, dithering introduces random noise in the signal, so that long sequences of smoothly varying samples are broken up, thus effectively avoiding the banding associated to such regions. Since then, dithering has become ubiquitous in all forms of digital signal coding, an enourmous body of work has been written on the subject, with several variants proposed (see [9] for a review for the case of digital images).

More closely related to our work, is the concept of *deterimistic dithering*, where the "noise" to be added is a function known both to the encoder and the decoder. This idea was first proposed in the context of sampling theory in [10], using deterministic pseudo-random noise sequences to contaminate the input signal. The contaminated signal is then sampled and quantized (non-predictively) to one bit per sample. Under certain conditions on the dithering sequence and the sampling rate, the method is shown to reconstruct a wide range of signals. This idea was later extended in [11] to non-pseudo-random dithering functions such as sinusoids, again in a non-predictive sampling context.

As with the preceding cases, our motivation lies in the removal of banding artifacts due to quantization. However, contrary to all of the above methods, we break the bands by allowing a pseudo-randomly chosen set of samples to be encoded *losslessly*. In this way, not only bands are removed, but the quantization error feedback that produces them is effectively broken often and, more importantly, at random positions, thus avoiding the formation of banding patterns typical of near-lossless predictive coding.

We apply this technique to a simplified version of the JPEG-LS standard [6], with clearly positive results in terms of overall mean squared error (MSE) and visual quality, both perceived subjectively, and as given by the Structural SIMiliarity (SSIM) image quality index [12], at the cost of a small (and predictable) increase in file size (measured in average bits per pixel – BPP). Moreover, the technique allows one to vary the amount of dithering, thus allowing the user to select different trade-offs between visual quality and compressibility.

2 Background

2.1 Predictive coding

Let $\mathbf{x}_1^n$ denote a sequence of n data samples to be encoded, where $\mathbf{x}_i^j$ is the subsequence from i to j; the sub-index may be omitted when $i = 1$. Coding of a new sample x_j is done by first computing a prediction of its value in terms of past samples, $\hat{x}_j = f(\mathbf{x}^{j-1})$, and then encoding (using some sort of Entropy coding) the prediction error $e_j = x_j - \hat{x}_j$. Since both the encoder and the decoder have access to the same information, the above procedure can be replicated at the decoder, so that only errors need to be transmitted.

Usual predictors include adaptive linear functions, $f(\mathbf{x}^{j-1}) = \sum_{k=1}^{p} a_k x_{j-k}$, and simple fixed predictors such as the constant ($\hat{x}_j = x_{j-1}$), and linear ($\hat{x}_j = 2x_{j-1} - x_{j-2}$) ones. The latter two are popular in "low complexity" compression algorithms such as JPEG-LS, as they require very little hardware resources. In order to compensate for the simplicity and fixed nature of these predictors, an adaptive component is usually included in the form of a *bias correction term* $b_j = \frac{1}{j-1}\sum_{i=0}^{j-1} e_i$. In this way, the final error $e_j = \hat{e}_j + b_j$, where $\hat{e}_j$ is the output of the fixed predictor, has an empirical distribution centered at 0, which results in compression gains. As fixed predictors tend to exhibit different biases depending on the local shape of the sequence, bias correction is often made *context dependent*, where by context we mean some function of the past few samples which captures the shape of the signal near the sample being encoded.

2.2 Near-lossless coding

In this setting, prediction errors are quantized in steps of size $\Delta = 2\delta + 1$, for a maximum absolute per-sample distortion of δ between the original signal $\mathbf{x}^n$ and the one reconstructed at the decoder, which we denote by $\mathbf{y}^n$. The quantized error $\tilde{e}_j$ is obtained from e_j via,

$$\tilde{e_j} = q(e_j) = \mathrm{sign}(e_j)\left[\frac{|e_j| + \delta}{1 + 2\delta}\right],$$

where $[\cdot]$ denotes rounding to nearest integer. However, since both decoder and encoder must have the same data available when processing the j-th sample, on both sides the prediction $\hat{x}_j$ must be now based on the reconstructed samples $\mathbf{y}^{n-1}$, $\hat{x}_j = f(\mathbf{y}^{j-1})$, and not on the original ones, $\mathbf{x}^{k-1}$. Therefore quantization

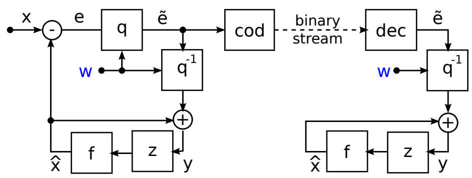

Fig. 2: Interleaved quantization encoding/decoding scheme. Here f stands for the predictor, q for the quantization block, q^{-1} for the inverse quantization block, $q^{-1}(\hat{e}_j) = 2\delta\hat{e}_j$, and z for a delay block, whose output is its input delayed by one time step for $j > 1$, and 0 for $j \leq 1$ (time indexes are ommited for simplicity).

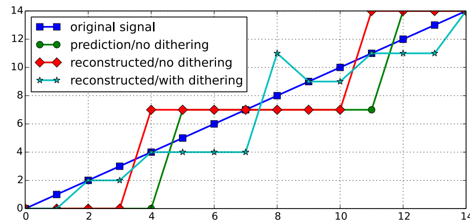

Fig. 3: Banding effect on lossy predictive coding of 1D signals, and the effect of dithering. Here $x_j = j-1$, $\hat{x}_j = x_{j-1}$, and $\delta = 3$. The last curve, shown in cyan, corresponds to a pseudo-random interleaved quantization of the prediction errors with a quantization probability of $p = 0.7$.

affects not only the transmitted errors, but also the prediction itself. In regions of the input signal where $|x_j - x_{j-1}| \ll \delta$ for many consecutive samples, the corresponding regions in the reconstructed signal $\mathbf{y}^n$ will be "flattened out", as small consecutive errors will be quantized to 0. To illustrate the above situation, consider the simple zero order predictor $\hat{x}_j = f(y^{j-1}) = y_{j-1}$. In this case, the unquantized error will be $e_j = x_j - \hat{x}_j = x_j - y_{j-1}$. Now, if $e_j = x_j - y_{j-1} < \delta$ we have that $\tilde{e}_j = 0$, in which case $y_j = y_{j-1}$. This error feedback loop goes on until $|\hat{x}_j - x_j| \geq \delta$, at which point a jump of size Δ will occur. This is illustrated in Figure 3, along with the proposed method, to be discussed next.

3 Interleaved quantization

We propose a simple modification to the lossy scheme presented in Section 2.2 where only a fraction $0 \leq p < 1$ of the prediction error samples are quantized. As can be seen in Figure 3 (cyan line), this is enough to break the staircase (1D banding) effect observed when no dithering is performed.

There are many possible ways to define the locations where quantization will occur. The algorithm that we present here, coined *interleaved quantization* (IQ) chooses such locations by generating a pseudo-random Bernoulli sequence $\mathbf{w}^n$ where $w_j \in \{0,1\}$, with $P(w_j = 1) = p$, and then quantizes the errors e_j at those locations j for which the corresponding $w_j = 1$. Although not truly random, the sequence $\mathbf{w}^n$ is sufficiently irregular to avoid generating visible artifacts in $\mathbf{y}^n$. *The key point here is that, given a fixed pseudo-random number generator, and a fixed seed, both the coder and the decoder know the exact places where quantization is, or is not, performed, without the need for encoding such places explicitly.* Other forms of interleaving are also possible. For example, the value of δ applied to each sample could be drawn uniformly between 0 and $\delta_{\max}$. A block diagram of the above procedure is presented in Figure 2.

A simplified analysis, which leaves the (positive) effect of "breaking the staircases" observed above aside, reveals that, for an IQ scheme with probability p, the output of the encoder can be seen as an interleaved coding of two sources: one corresponding to a lossy signal, and other to a lossless one. Therefore, if L_{lossy} is the codelength obtained for a given case with the fully lossy scheme ($p = 1$), and L_{lossless} is the one obtained in the lossless case, the resulting code length for the IQ scheme L_{iq} should be close to $pL_{\text{lossy}} + (1-p)L_{\text{lossless}}$. Also ignoring the "staircase breaking effect", and with similar arguments, the distortion D_{iq} in the image reconstructed by IQ should be close to $pD_{\text{lossy}} + (1-p)D_{\text{lossless}}$. As will be shown in Section 4, these simplified results are indeed quite accurate. In this way, p serves as an additional parameter to select a particular rate-distortion trade-off.

We applied the IQ idea to a simplified version of JPEG-LS which we will refer to as "IQ" in the sequel. Its main difference with JPEG-LS lies in the way that it switches to *run coding mode*, which in IQ is analogous to the lossless case, whereas JPEG-LS takes into account quantization (a reasonably complete description of JPEG-LS is not possible given the space constraints; please refer to [6] for technical details on its definition). As JPEG-LS uses a very simple, fixed (2D) predictor together with a context-dependent bias correction term, the effect of quantization fits well within the simplified analysis of Figure 3, as the results below show.

4 Results and discussion

The primary purpose of our algorithm is to improve upon the visual artifacts produced by current near-lossless prediction-based image coding techniques. Figures 4 and Figure 5 are examples for which such artifacts are clearly visible, even for small target distortions, on a very low dynamic range medium such as paper or even an ordinary computer monitor. It is important to underline that such effects are much more noticeable, at even smaller target distortions, on current commercial displays aimed at consumers in general. For a numerical evaluation of our method, we applied the IQ algorithm to a grayscale version of the "Kodak dataset" [1]; [2]. In Figure 6 we report these results in terms of the traditional Rate-Distortion curve, based on mean squared error (MSE), and on a Rate-Quality curve, with the "quality" given by the Structural SIMiliary index [12]. In both cases, we compare our results against the classic (lossy, and not predictive, but transform-based) JPEG [13], and JPEG-LS [6], focusing on the near-lossless region ($\delta \leq 5$).

As can be seen in Figure 6(left), from a traditional quadratic Rate-Distortion perspective, the proposed interleaved quantization does not offer any advantages over JPEG-LS; it essentially coincides coincides (as expected) with JPEG-LS for $p = 1$, and moves upwards (this is worse) as p decreases. Also as expected, both

[1] Publicly available at `http://r0k.us/graphics/kodak/`

[2] Additional examples, as well as the source code, are available at `http://iie.fing.edu.uy/~nacho/demos/iq/`.

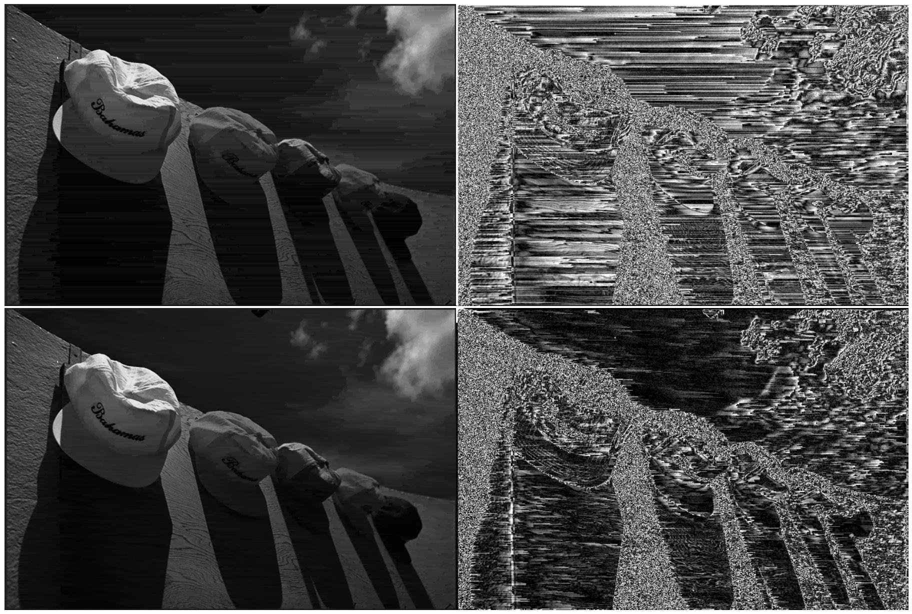

Fig. 4: Sample grayscale results. Here we show a grayscale version of the "kodim03" image from the Kodak dataset. The above pictures correspond to the near-lossless compression of kodim03 for $\delta = 10$ and no interleaved quantization $p = 1.0$ (0.86 bpp), and its absolute error with respect to the original undistorted image. The bottom row shows the same image, and its error, when compressed using interleaved quantization with $p = 0.9$ (0.96 bpp). In this case, the artifacts are dramatically reduced at a slight bitrate increase of 0.1 bpp.

give better R-D tradeoffs than JPEG in the low-distortion region shown. In terms of the Rate-SSIM curve shown in Figure 6(right), however, IQ improves over JPEG-LS in the very low distortion region, with several configurations lying below and to the right of the JPEG-LS curve (that is, higher SSIM at the same bitrate). At some point (here, below 3.00bpp), also as expected, both IQ and JPEG-LS lose by a significant margin to the classic JPEG algorithm, which is optimized for non-near lossy operation. Although these numerical results may seem dissapointing, we argue that the gain in terms of visual quality, as evidenced in figures 4 and 5, is much larger. Also, it is important to bear in mind that one of the advantages of near-lossless compression lies in its guaranteed maximum distortion, something which may be advantageous, from a legal standpoint, over traditional methods such as JPEG (example, medical imaging for diagnosis). In this sense, our method retains such advantages, while producing less visual artifacts, and at a small decrease in compression rate.

Fig. 5: Color example. Without quantization (left column) the banding effect is already noticeable for $\delta = 5$, is clearly visible for $\delta = 10$. Leaving only 10% unquantized already improves the visual quality significantly, as can be seen on the right column. *Note: this example is best appreciated on a computer screen.*

References

1. Elias, P.: Predictive coding–I. Information Theory, IRE Transactions on **1**(1) (1955) 16–24
2. Elias, P.: Predictive coding–II. Information Theory, IRE Transactions on **1**(1) (1955) 24–33
3. O'Neal, J.: Predictive quantizing systems (differential pulse code modulation) for the transmission of television signals. Bell System Technical Journal, The **45**(5) (1966) 689–721
4. Golomb, S.W.: Run-length encodings (corresp.). Information Theory, IEEE Transactions on **12**(3) (1966) 399–401
5. Rice, R.F.: Some practical universal noiseless coding techniques–part III. Technical Report JPL-91-3, Jet Propulsion Laboratory (1991)
6. Weinberger, M., Seroussi, G., Sapiro, G.: The LOCO-I lossless image compression algorithm: Principles and standardization into JPEG-LS. IEEE Trans. IP **9** (2000)
7. Cutler, C.: Differential quantization of communication signals (1952) US Patent 2,605,361.
8. Roberts, L.G.: Picture coding using pseudo-random noise. IRE Transactions on Information Theory **8**(2) (1962) 145–154
9. Chen, J.S.: A comparative study of digital halftoning techniques. In: Aerospace and Electronics Conference, 1992. NAECON 1992., Proceedings of the IEEE 1992 National. (1992) 1139–1145 vol.3
10. Masry, E.: The reconstruction of analog signals from the sign of their noisy samples. IEEE Trans. IT **IT-27**(6) (1981) 735–745

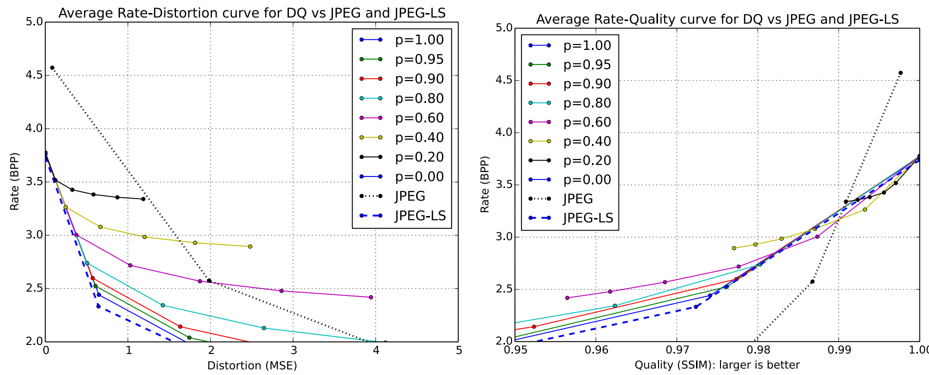

Fig. 6: Summary of results. Left: Rate-Distortion curves in terms of MSE for JPEG, JPEG-LS and IQ in the near-lossless region. Right: Rate-SSIM curves for the same algorithm. In both cases, the IQ curve is plotted for various values of the quantization probability p, each dot corresponding to a value of δ, starting with $\delta = 0$ (lossless, that is, MSE$= 0$ on the left, and SSIM$= 1.0$ on the right), in increments of 1. In the first case, the Rate-Distortion curves for IQ are above the JPEG-LS one for $p < 1.0$, meaning that the process of interleaving does not provide advantages over JPEG-LS in terms of quadratic Rate-Distortion. (It remains, however, below that of JPEG, for values of $\delta < 4$.) In the second case, where lower-right is better, IQ improves over JPEG-LS in the very low distortion region ($\delta \leq 3$).*Note: this is a color graph.*

11. Cvetković, Z., Daubechies, I.: Single bit oversampled A/D conversion with exponential accuracy in the bit rate. In: Proc. Data Compression Conference (DCC) 2000, Snowbird, UT. (2000) 343–352
12. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error measurement to structural similarity. IEEE Trans. IP **13**(1) (2004)
13. Wallace, G.: The JPEG still picture compression standard. Communications of the ACM **34, 4** (1991) 30–44