Rate-Distortion Optimal Adaptive Quantization and Coefficient Thresholding for MPEG Coding

Dirk Farin[1], Michael Käsemann,
Peter H.N. de With[2], and Wolfgang Effelsberg[1]

[1] Dept. of Computer Science IV, University Mannheim, Germany
farin@informatik.uni-mannheim.de
[2] CMG Eindhoven / University of Technology Eindhoven, Netherlands
p.h.n.de.with@tue.nl

Abstract. This paper presents an MPEG-2 compatible adaptive quantization algorithm that leads to the optimal encoding of I-frames in the sense of maximizing PSNR. It integrates three key features into a single Lagrangian optimization model: adaptive quantization including quantizer-change overhead consideration, coefficient thresholding, and a new coefficient amplitude reduction technique. Our results show that I-frames generated by the TM5 reference encoder are about 1.5-2.0 dB below the theoretical optimum.

1 Introduction

In video coding, it is desired to obtain the best possible image quality at a predefined bit rate. Typically, MPEG encoders control the quantization step size based on a feed-forward control algorithm. This results in frequent changes of the quantizer setting, thereby generating a coding overhead due to the coding of new quantizer settings. Moreover, in single-pass video coders (like TM5 [6]), a bad prediction of image coding complexity may lead to an unequal quality distribution in a single image.

The principle of optimal quantizer selection has been introduced in [5], where the Lagrange-multiplier method was used to assign quantizers to independent coding units to optimize the overall quality under a limited bit-budget. However, changing quantizer settings in the MPEG standard involves a coding overhead. Hence, a quantization scale cannot be chosen independent of the context. In [3], an extension to the quantizer selection algorithm is proposed which takes quantizer-change overhead into account. A different approach to increase coding efficiency has been proposed in [4] (coefficient thresholding). The algorithm omits DCT coefficients when the amount of bits saved is high but the decrease of image quality is relatively small.

This paper combines the above-mentioned techniques into a unified optimization process. Furthermore, we implemented *coefficient amplitude reduction* (CAR) as a new technique for improving image quality. It is based on the fact that reducing the amplitude of a DCT coefficient can be advantageous if the reduced coefficient has considerably shorter Huffman codes (consider lengthy escape codes that are reduced to short code-book entries).

2 Adaptive Quantization

Consider the problem of coding an image at rate[3] R_{max} with minimum distortion (MSE) D. Each image consists of a fixed number of coding units (e.g., macroblocks), which can each be coded with different quantizer settings q_i. Let $D_i(q_i)$ be the distortion of macroblock i when quantized with q_i, and let $R_i(q_i)$ be the number of bits required for coding the macroblock. The optimization problem can now be formulated as

$$\min_{q_i} \sum_i D_i(q_i) \quad \text{such that} \quad \sum_i R_i(q_i) \leq R_{max}. \tag{1}$$

In [5], Shoham showed that by using the Lagrange-multiplier framework, this constrained optimization problem can be written as the equivalent problem

$$\min_{q_i} \sum_i D_i(q_i) + \lambda R_i(q_i) \tag{2}$$

for a fixed λ. The paper [5] also provides the proof that each solution of the transformed (unconstrained) problem is also a solution of the original problem with the rate-constraint $R'_{max} = \sum_i R_i(q_i)$ if the rate-distortion function is convex. As R'_{max} is dependent on λ, a suitable value of λ has to be determined to solve the original problem with $R'_{max} = R_{max}$ [4]. The suitable λ-value can be determined using a binary search.

The advantage of the second problem formulation in eq. (2) is that the sum and the minimum operator can be exchanged to

$$\sum_i \min_{q_i} D_i(q_i) + \lambda R_i(q_i). \tag{3}$$

This formulation obviously reveals that the global optimization can now be carried out independently for each macroblock, making an efficient implementation feasible.

Unfortunately, according to the MPEG standard, changing the quantization scale requires additional bits in the macroblock header to code the new settings. The overhead comprises 2 bits for coding the macroblock mode and 5 bits for the quantization scale, compared to only 1 bit for the macroblock mode when the quantizer is the same as in the last macroblock. Especially at low bit rates, this overhead cannot be ignored. Hence, we introduce the quantizer change overhead as an extra contribution to the rate:[5]

$$R^{ov}(q_i, q_{i-1}) = \begin{cases} 1 & \text{for} \quad q_i = q_{i-1}, \\ 7 & \text{for} \quad q_i \neq q_{i-1}, \\ 1 & \text{if } q_i \text{ is the first MB in a slice.} \end{cases} \tag{4}$$

[3] In this paper, we use the term *rate* to denote the number of bits per frame.

[4] In practice, exact equivalence of rate cannot be guaranteed, and a suitable tolerance has to be accepted.

[5] Note that additional header fields exist. As they have constant size, they can be ignored in the minimization problem. However, they have to be considered when calculating the total rate.

After adding the overhead to the functional in (2), our optimization problem reads as

$$\min_{q_i} \sum_i D_i(q_i) + \lambda R_i(q_i) + \lambda R^{ov}(q_i, q_{i-1}). \tag{5}$$

This can no longer be solved independently for each macroblock, but can be determined using a dynamic programming approach. Consider the graph in Figure 1, in which each column of nodes represents a macroblock and each row defines a quantizer scale. Each path through the graph corresponds to a possible coding of the frame. Traversing the nodes induces associated costs $D_i(q)+\lambda R_i(q)$ and graph edges from row q_1 to row q_2 have costs $R^{ov}(q_2, q_1)$. Hence, the total path cost is equivalent to functional (5), and the minimum cost path defines the solution of the above minimization problem.

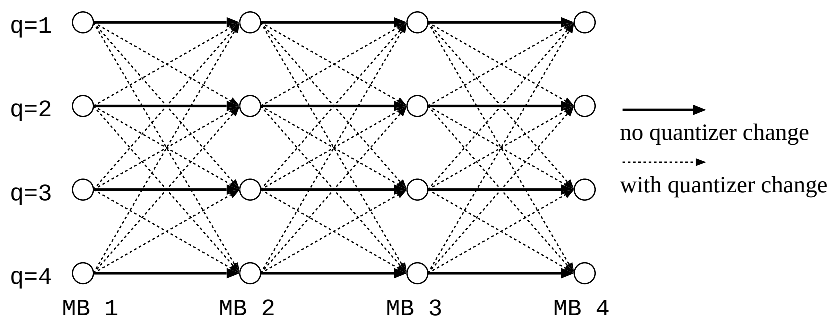

Fig. 1. Equivalent graph search problem to the Lagrangian minimization problem. Note that there are 31 different quantizer scales in MPEG instead of only four.

3 Thresholding

Ramchandran [4] introduced coefficient thresholding as a post-processing step after quantization to further reduce the bit rate while still retaining as much image quality as possible. The idea is to drop coefficients (set to zero) when the additional distortion is small compared to the number of bits saved. Considering thresholding as a separate post-processing step induces the difficulty that it is not clear how to choose quantization parameters. If a target rate of R_{max} is requested, obviously the rate after quantization has to be greater, so that thresholding can be used to further decreasing the rate. However, the exact rate is unknown. We solved this problem by incorporating coefficient thresholding together with adaptive quantization into a single Lagrangian framework.

The following algorithm exploits a useful property of the DCT which leads to an efficient implementation of thresholding. As the DCT does not change the ℓ_2 norm of a vector, the MSE of a block can either be computed in the spatial domain or, equivalently, in the frequency domain. This property enables to calculate efficiently the additional distortion that is introduced by modifying (or even omitting) a single coefficient.

To simplify notation, we concentrate on a single DCT block, consisting of coefficients c_i. We denote quantization by $c'_i = Q(c_i)$ and dequantization by $\hat{c}_i = Q^{-1}(c'_i)$. Let $C = \{(p_i, \hat{c}_i)\}$ be the ordered set (ascending p_i) of non-zero quantized coefficients ($\hat{c}_i \neq 0$), where p_i is the position of the coefficient (in zig-zag order). Hence, by using a table of the Huffman code-lengths $r(run, value)$, the bits needed to code coefficient i are $r(p_i - p_{i-1} + 1, c'_i)$. Omitting the coefficient would induce additional distortion c_i^2. Note that coefficient $i = 0$ is always the DC coefficient which cannot be omitted. Let $S \subseteq C$ be the subset of coefficients in the block that we decide to code. Hence, we intend to minimize the Lagrangian cost associated to a selection S:

$$\min_{S \subseteq C} \underbrace{\sum_{(p,c) \in C-S} c^2}_{\text{costs for omitting}} + \underbrace{\sum_{1 \leq i \leq |S|} (c_i - \hat{c}_i)^2 + \lambda\, r(p_i - p_{i-1} + 1, c'_i)}_{\text{costs for coding quantized coefficient}} . \quad (6)$$

Similar to adaptive quantization in the previous section, this minimization problem can be solved by computing an equivalent graph search. The corresponding graph is depicted in Figure 2. Every non-zero coefficient is represented by a graph node. A special node *EOB* is added as a last node so that skipping the last coefficient is possible. Each non-skipping edge is attributed with weight $(c_i - \hat{c}_i)^2 + \lambda\, r(p_i - p_{i-1} + 1, c'_i)$, consisting of the quantization error and the length of the Huffman code. Each skipping edge is attributed with weight $\sum c_i^2$, where the sum includes all skipped coefficients.

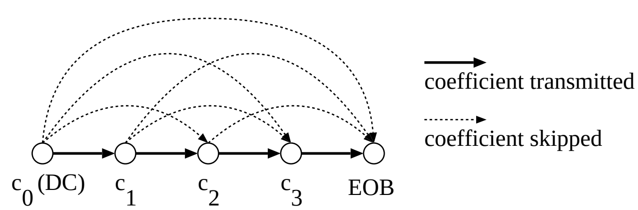

Fig. 2. Equivalent graph search problem to coefficient thresholding.

To visualize the principle of coefficient thresholding, we coded a frame using fixed quantization scales. Afterwards, we applied coefficient thresholding to further reduce the bit rate. The result for a frame of the *Claire* sequence is shown in Fig. 5. It can be seen that for small reductions of rate in the thresholding step, the slope of the rate-distortion curve is less than that using quantization only. However, for larger rate reductions, the slope of the thresholding curves becomes much steeper, corresponding to a faster decrease of image quality.

To optimally join adaptive quantization and thresholding, we merged the adaptive quantization graph and the thresholding graph into a single combined graph (Fig. 3). In this way, we get the "convex hull" over the curves of Fig. 5, being the optimal combination of adaptive quantization and thresholding at every bit rate.

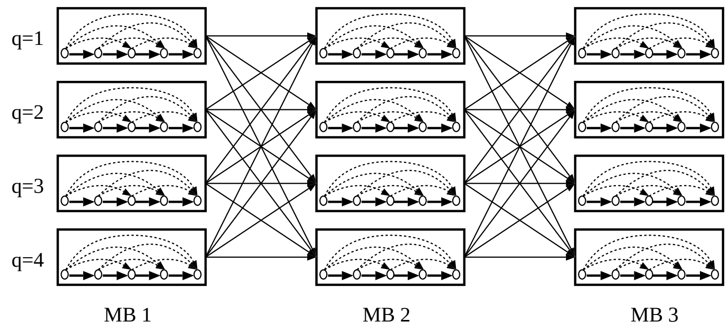

Fig. 3. Combination of adaptive quantization graph and thresholding graph. Each box shown with a thresholding graph actually contains six concatenated thresholding graphs (for the six DCT blocks contained in each macroblock).

4 Coefficient Amplitude Reduction

In this section, we introduce *coefficient amplitude reduction* (CAR) as a generalization of coefficient thresholding. The idea is that it can be advantageous to decrease the value of a coefficient when the number of bits saved outweighs the additional distortion. Especially when the true coefficient value is near the lower decision boundary of the quantization interval, reducing the coefficient amplitude does not introduce much additional distortion (Fig. 4a). On the other hand, when a slight decrease prevents the run-value pair to be coded with costly escape sequences, the bit rate gain may be significant. As the MPEG Huffman table is monotone, increasing coefficients will never lead to shorter codes.

CAR can be implemented by extending the thresholding graph as shown in Fig. 4b. For each coefficient with value c'_i, further c'_i nodes are created, representing the new (reduced) value of the coefficient (range $1, \ldots, c'_i$). Node costs are assigned accordingly. All new edges and dummy nodes have zero cost.

CAR can be implemented independently of thresholding by omitting the skipping edges in Fig. 4b. A further advantageous property is that the combination of CAR and thresholding can be implemented by successive application of CAR first and thresholding afterwards on the modified coefficients. Clearly, the CAR step selects the optimal coefficient value, equivalent to the shortest path between every second node. As every sub-path of a shortest path is also a shortest path, either this path or a skipping edge would be chosen by the thresholding step. By replacing the thresholding sub-graphs in Fig. 3 with the CAR graphs, we obtain the theoretically optimal encoding of the frame.

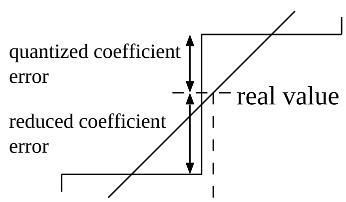

(a) Quantization error does not increase much when original value is near decision boundary.

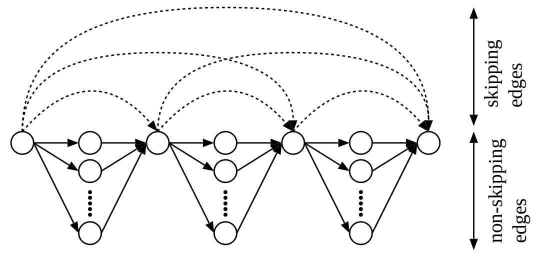

(b) Coefficient thresholding graph, enlarged with additional amplitude reduction nodes.

Fig. 4. Coefficient amplitude reduction.

5 Results

We have implemented the above algorithms into the SAMPEG encoder framework [1]/[2]. A single frame was selected from a test sequence and coded with different combinations of quantization and coefficient modification. For quantization, we used three variants: constant quantization (all macroblocks are coded with the same quantization scale), adaptive quantization (as explained above), and adaptive quantization without considering quantizer-scale change overhead (NOO). Furthermore, we used the TM5 reference implementation [6] for comparison. For coefficient modification, we used: thresholding alone, CAR alone, both combined, and both disabled.

Table 1 shows the absolute PSNR reached for a fixed rate and the increase of PSNR compared to using constant quantizer scales. Adaptive quantization increases the PSNR by about 0.13 dB. Applying thresholding leads to another 0.2-0.3 dB increase. The PSNR increase obtained from CAR is only marginal and can be neglected. NOO cannot increase PSNR much above the constant quantization heuristic. At low bit rates, it even performs worse because of frequent quantizer changes. Comparable results are obtained for other input images.

According to our results, using frame-constant quantization scales is a good heuristic for PSNR optimal quantization. The heuristic achieves $\approx +1.2-1.8$ dB compared to TM5 and is only $\approx 0.3-0.5$ dB below the theoretical maximum.

In a second experiment, we examined which coefficients in a DCT block were thresholded most. Approximately 70-80% of the thresholded coefficients were at the end of the DCT block. Skipping these coefficients moves the EOB code to an earlier position, which results in a particulary large reduction of bits. Accordingly, applying thresholding to only the last coefficients of a block results in about 80% of the PSNR increase. This fact may enable computationally inexpensive heuristics for fast thresholding.

Claire (≈67800 bits)	no coeff. mod.	CAR	thresholding	both
constant q.scale[6]	46.16 (0.0)	n/a	n/a	n/a
TM5	44.91 (−1.25)	n/a	n/a	n/a
adapt./no overh. (NOO)	46.21 (+0.06)	46.22 (+0.07)	46.40 (+0.24)	46.41 (+0.25)
adaptive quant.	46.29 (+0.13)	46.30 (+0.15)	46.46 (+0.30)	46.47 (+0.31)

Stefan (≈170000 bits)	no coeff. mod.	CAR	thresholding	both
constant q.scale	36.87 (0.0)	n/a	n/a	n/a
TM5	35.62 (−1.25)	n/a	n/a	n/a
adapt./no overh. (NOO)	36.96 (+0.09)	36.96 (+0.10)	37.27 (+0.40)	37.27 (+0.40)
adaptive quant.	36.99 (+0.12)	37.00 (+0.13)	37.29 (+0.42)	37.30 (+0.43)

Table 1. Overall results: PSNR in dB (absolute and increase compared to constant quantization). The bits per frame were chosen as if the sequence was coded at 1.2 Mbps (at CIF resolution).

6 Conclusions

We have presented an optimal quantization algorithm for MPEG coded I-frames. It achieves to generate images with the best possible image quality at a given bit rate. Even though it may be computationally too complex for practical encoding applications, it is suitable to serve as a reference to compare other, heuristic, algorithms with. The algorithm can easily be extended to support P- and B-frames by including the additional macroblock mode decisions in a similar way.

References

1. Dirk Farin, Niels Mache, and Peter H. N. de With. SAMPEG, a scene adaptive, parallel MPEG-2 software encoder. *SPIE Visual Communications and Image Processing*, 4310:272–283, January 2001.
2. Dirk Farin, Niels Mache, and Peter H. N. de With. A software-based high-quality MPEG-2 encoder employing scene change detection and adaptive quantization. *IEEE ICCE Digest*, pages 148–149, June 2001.
3. Antonio Ortega and Kannan Ramchandran. Forward-adaptive quantization with optimal overhead cost for image and video coding with applications to MPEG video coders. In *SPIE Digital Video Compression*, February 1995.
4. Kannan Ramchandran and Martin Vetterli. Rate-distortion optimal fast thresholding with complete JPEG/MPEG decoder compatibility. *IEEE Transactions on Image Processing*, 3(5):700–704, September 1994.
5. Yair Shoham and Allen Gersho. Efficient bit allocation for an arbitrary set of quantizers. *IEEE Transactions on Acoustics and Signal Processing*, 36(9):1445–1453, September 1988.
6. MPEG-2, Test Model 5 (TM5). *Doc ISO/IEC JTC1/SC29/WG11/93-225b*. Test Model Editing Committee, April 1993.

[6] Even though thresholding can be combined with constant quantization, the performance highly depends on the selected rate (see Fig. 5).

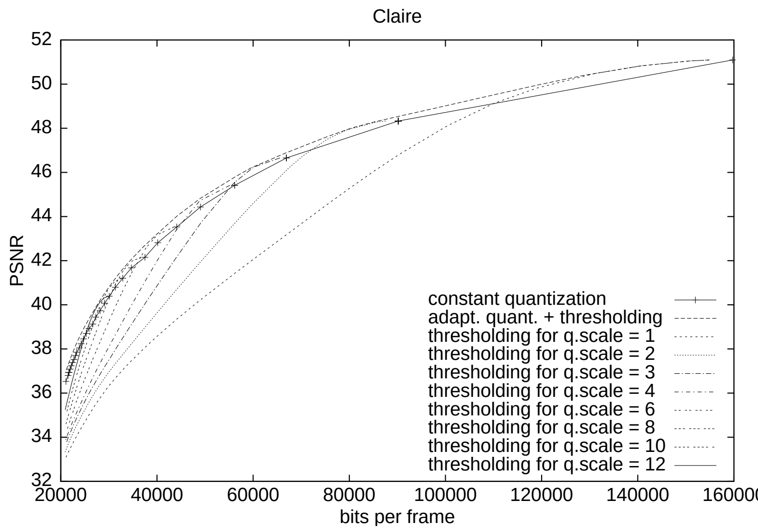

Fig. 5. Using coefficient thresholding as independent post-processing step after quantization.

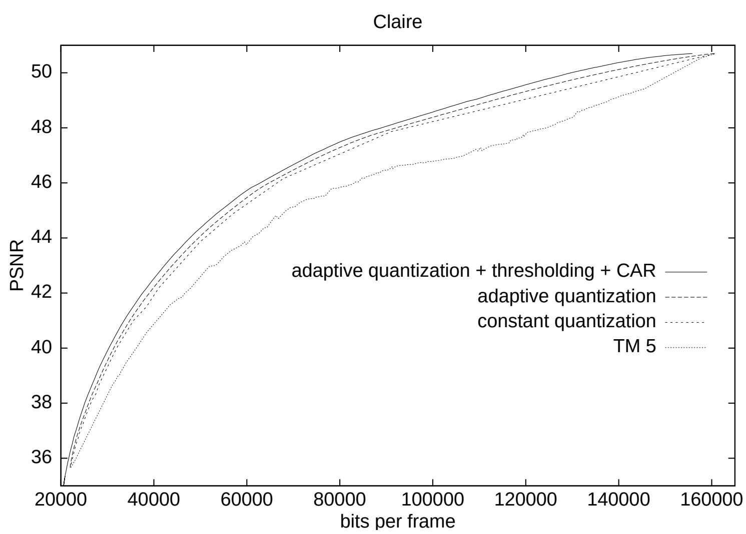

Fig. 6. Results for the *Claire* sequence.