

Title: Transform Skipping in the presence of Scaling Lists

Status: Input Document to JCT-VC

Purpose: Proposal

Author(s): David Flynn
Andrea Gabriellini
Marta Mrak
Matteo Naccari

davidf@rd.bbc.co.uk
andrea.gabriellini@bbc.co.uk
marta.mrak@bbc.co.uk
matteo.naccari@bbc.co.uk

Source: BBC R&D

Abstract

The 9th meeting of the JCT-VC adopted transform skipping (JCTVC-I0408[1]) into the working draft text. Transform skipping allows 4x4 TUs to bypass the transform and instead directly quantise the prediction residue in the spatial domain. An observation was made that there is a possible undesired interaction with the adopted proposal in the presence of a non-flat scaling list, whereby the frequency weighted scaling list is applied to spatial coefficients. This contribution examines the interaction and proposes a method to derive a flat scaling list for use by such TUs.

Analysis

The quantification of any detriment to coding efficiency caused by the use of a non-flat quantisation matrix is difficult. The use of scaling lists is known to increase the distortion as seen by the objective measurement of PSNR (which is similarly true for an increase in QP). It naturally follows that using a flat quantisation matrix or lower QP for a single TU size will lead to an improvement in the objective measure. As such, measurements that show BD-rate improvements are not particularly helpful (in the extreme, they would lead to the removal of the scaling lists).

Two obvious questions to consider are, what is the effect of using a non-flat quantisation matrix on TUs with skipped transforms, and what happens if a flat matrix is used instead.

The analysis presented here is performed using the default scaling lists for 4x4 TUs. For this size of TU it is possible to use six different scaling lists (for intra and inter, each with three colour components). The default scaling list is only differentiated by intra and inter modes (each colour component is otherwise identical).

The effect of a non-flat quantisation matrix

The magnitude of the default quantisation matrix for frequency domain coefficients increases towards the bottom row and right-hand column, as shown in Figure 1. The effect of such a matrix is to increase QP for some coefficients by a range effectively up to QP+12¹

Figure 1: Default quantisation matrix $m_{i,j}$ for 4x4 TUs and all colour components

$$\begin{array}{cc} \left[\begin{array}{cccc} 16 & 16 & 17 & 21 \\ 16 & 17 & 20 & 25 \\ 17 & 20 & 30 & 41 \\ 21 & 25 & 41 & 70 \end{array} \right] & \left[\begin{array}{cccc} 16 & 16 & 17 & 21 \\ 16 & 17 & 21 & 34 \\ 17 & 21 & 24 & 36 \\ 21 & 24 & 36 & 57 \end{array} \right] \\ \text{(a) Intra} & \text{(b) Inter} \end{array}$$

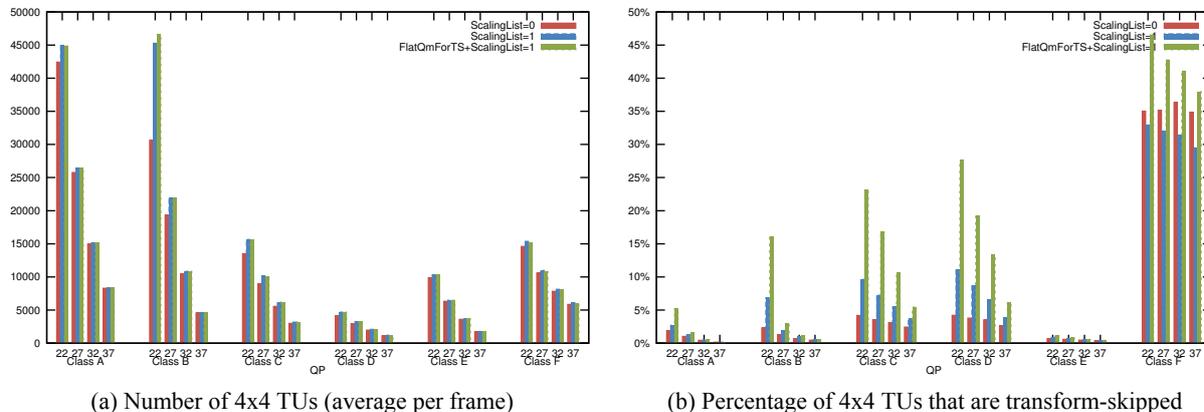
When this matrix is applied to the residual of a transform-skipped TU, the effect is to quantise the bottom and right-hand edges of the TU more heavily than the top and left. For a subsequent neighbouring intra

¹lower-right coefficient for Intra, $\log_2(70/16) * 6 = 12.76$

PU, this can potentially have an effect on the quality of the predictor used for that PU.

By using the default scaling list (with otherwise the HM-7.0 common conditions), in classes A–E there is an increase in both the number of 4x4 TUs (Figure 2a) and the percentage of these TUs that are transform-skipped (Figure 2b). A notable exception are the class F sequences, where although there is a small increase in 4x4 TUs, there is a significant decrease in the incidence of transform-skipped TUs.

Figure 2: Effect of flat quantisation matrix on block selection i_he10 configuration



Another interesting effect is the change in distribution of significant coefficients within the 4x4 transform-skipped TUs for class F. Ordinarily, without the default quantisation matrix, the incidence of significant coefficients increases towards the bottom right of the TU. With the default scaling list, the incidence decreases towards the final coefficient.

To determine if there is any change in subjective quality when transform skipping is used in the presence of scaling lists, an informal visual comparison was conducted using:

- Common conditions[2] with scaling lists enabled and intra transform skip disabled for all configurations ($-ScalingList=1, -TS=0$).
- Common conditions with scaling lists enabled and intra transform skip enabled for all configurations ($-ScalingList=1, -TS=1$).

When visually inspecting sequences with a high incidence of transform-skipped TUs, no obvious differences were found.

The effect of forcing a flat quantisation matrix

If we consider a case where the transform-skipped 4x4 TUs are forced to use a flat quantisation matrix (without affecting non-skipped TUs), a few predictions can be made. That transform-skipped TUs will be quantised less, allowing them to more faithfully reproduce the residual, thus improving the PSNR. Less quantisation will result in an increase in bitrate. It is likely that more 4x4 blocks will use the transform-skipped mode due to the significant improvement in distortion.

To test this, an experiment is performed using HM-7.0 for intra configurations and HM-7.0 with the proposed modifications for Inter Transform Skipping[3] for non-intra configurations using the common conditions with scaling lists enabled and intra transform skip enabled for all configurations ($-ScalingList=1, -TS=1$). The behaviour is compared to a modified encoder that derives a flat quantisation matrix for use by transform-skipped TUs. Due to the aforementioned reasons, BD-Rate results are not indicative of actual rate savings, but are nonetheless presented in Table 1.

The act of then forcing the 4x4 transform-skipped TUs to use a flat quantisation matrix results in a minor change in 4x4 mode utilisation (relative to $ScalingList=1$, Figure 2a), and a more significant increase in the proportion of transform-skipped TUs (Figure 2b). Examining the distribution (Figure 4) of the number of significant coefficients produced by transform-skipped TUs and the distribution of coefficient position within the blocks shows not only the propensity to increase the number of 4x4 transform-skipped

TUs, but in the cases of classes A–D, a change to the distribution of the number of significant coefficients reveals how the mode is being used in preference to the alternative candidates.

One possibility for this change in distribution is the aforementioned bias caused to the RDO search. By using a flat quantisation matrix and sending a large number of non-zero coefficients, such a large decrease in distortion is achieved that the increase in rate is considered worthwhile with the current lambda. The fact that there is an increase in transform skipping when ScalingLists are enabled, suggests again that either the distortion metric or lambda values are suboptimal when considering this mode.

An additional examination of the decoded pictures found no observable differences.

Possible solutions

Two possible solutions to the problem (each with two variants) are:

- 1.a Do nothing, allow encoders to use a non-flat matrix for 4x4 transform-skipped TUs.
- 1.b Do nothing, allow encoders to set a flat matrix for all 4x4 TUs.
- 2.a Force the quantisation matrix to be flat using the default value of 16 for affected 4x4 TUs. This would maintain the behaviour as if scaling lists were disabled.
- 2.b Require the decoder to derive a flat matrix from the default or signalled scaling list for affected 4x4 TUs.

Do nothing

Option one does not require any further work, and the true impact is uncertain. If larger block sizes were used this may be unattractive.

It certainly feels counterintuitive that it would be desirable to more heavily quantise the coefficients at the bottom and right edges of a TU that are subsequently used for intra prediction.

A mechanism to derive a flat matrix

Option two requires a minor alteration to the draft text. It is suggested to use option (2.b), even though the effect is the same as (2.a) in the presence of default scaling lists, in the case where the scaling list uses a value different from 16 for 4x4 TUs, option 2.b will not cause a disparity in quantisation between TUs with and without skipped transforms.

Figure 3: Proposed modification to draft text

8.6.3 Scaling process for transform coefficients

Inputs of this process are:

...
 – a variable `transSkipFlag` specifying whether transform is applied to the current block.
 ...

For the derivation of the scaled transform coefficients d_{ij} with $i = 0..nW - 1$, $j = 0..nH - 1$ the following applies

- The scaling factor m_{ij} is derived as follows.
 - If `scaling_list_enable_flag` is equal to 0,

$$m_{ij} = 16$$
 - If `scaling_list_enable_flag` is equal to 1 and `transSkipFlag` is equal to 1,

$$m_{ij} = \text{ScalingFactor}[\text{SizeID}][\text{RefMatrixID}][\text{trafoType}[0]]$$
 - Otherwise (`scaling_list_enable_flag` is equal to 1 and `transSkipFlag` is equal to 0),

$$m_{ij} = \text{ScalingFactor}[\text{SizeID}][\text{RefMatrixID}][\text{trafoType}[i \times nW + j]]$$

The proposed modification is illustrated in Figure 3. This can be implemented in a decoder by either making a decision for each transform coefficient (may be preferable for hardware) or by pre-computing an additional matrix m_{ij} for transform-skipped TUs (may be preferable for software).

For simplicity, the modifications made to HM use the first method, which introduces a per coefficient branch for all transform sizes in the quantiser, resulting in a measurable increase in runtime. No sane practical implementation would wish to do this.

Table 1: Comparison of HM-7.0 with default and flat quantisation matrices for transform-skipped TUs

(a) All Intra Main				(b) All Intra HE10			
	Y' BD-rate	U BD-rate	V BD-rate		Y' BD-rate	U BD-rate	V BD-rate
Class A	0.0%	-0.1%	-0.1%	Class A	0.0%	0.0%	0.0%
Class B	-0.1%	0.6%	0.6%	Class B	-0.1%	0.6%	0.6%
Class C	-0.5%	3.0%	3.2%	Class C	-0.3%	3.2%	3.2%
Class D	-0.5%	4.0%	3.8%	Class D	-0.3%	3.9%	3.8%
Class E	0.0%	0.0%	0.0%	Class E	0.0%	0.0%	-0.1%
Class F	-2.8%	1.0%	-0.7%	Class F	-2.6%	2.2%	0.6%
All (A-E)	-0.2%	1.5%	1.5%	All (A-E)	-0.1%	1.6%	1.6%
Enc Time	103%			Enc Time	103%		
Dec Time	100%			Dec Time	107%		

(c) Random Access Main				(d) Random Access HE10			
	Y' BD-rate	U BD-rate	V BD-rate		Y' BD-rate	U BD-rate	V BD-rate
Class A	0.1%	-0.5%	-0.4%	Class A	0.1%	-0.1%	-0.3%
Class B	0.1%	1.4%	1.4%	Class B	0.1%	1.1%	1.2%
Class C	-0.8%	1.0%	1.3%	Class C	-0.7%	1.2%	1.4%
Class D	-0.6%	3.1%	2.9%	Class D	-0.7%	3.0%	2.6%
Class E				Class E			
Class F	-3.8%	-2.5%	-3.8%	Class F	-3.7%	-1.3%	-2.5%
All (A-E)	-0.3%	1.3%	1.3%	All (A-E)	-0.3%	1.3%	1.2%
Enc Time	102%			Enc Time	101%		
Dec Time	105%			Dec Time	102%		

(e) Low delay B Main				(f) Low delay B HE10			
	Y' BD-rate	U BD-rate	V BD-rate		Y' BD-rate	U BD-rate	V BD-rate
Class A				Class A			
Class B	0.0%	1.5%	1.7%	Class B	0.0%	1.3%	1.2%
Class C	-0.3%	1.7%	1.7%	Class C	-0.2%	1.4%	1.5%
Class D	-0.2%	3.2%	3.4%	Class D	-0.2%	3.3%	3.0%
Class E	-0.1%	0.4%	0.2%	Class E	0.0%	-0.5%	-0.4%
Class F	-2.9%	-0.5%	-2.1%	Class F	-2.6%	0.1%	-1.8%
All (A-E)	-0.1%	1.8%	1.9%	All (A-E)	-0.1%	1.5%	1.4%
Enc Time	101%			Enc Time	101%		
Dec Time	103%			Dec Time	102%		

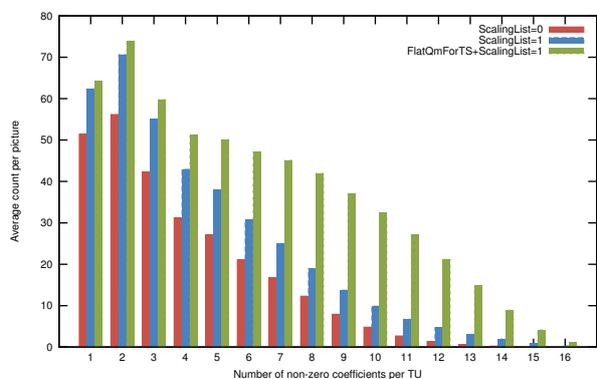
IPR declaration

The British Broadcasting Corporation may have current or pending patent rights relating to the technology described in this contribution and, conditioned on reciprocity, is prepared to grant licenses under reasonable and non-discriminatory terms as necessary for implementation of the resulting ITU-T Recommendation | ISO/IEC International Standard (per box 2 of the ITU-T/ITU-R/ISO/IEC patent statement and licensing declaration form).

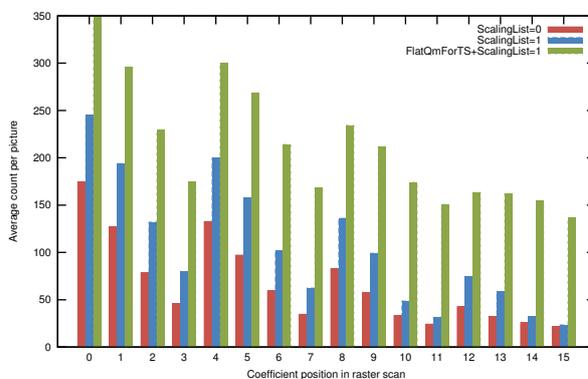
References

- [1] C. Lan, J. Xu, G. J. Sullivan, and F. Wu, "Intra transform skipping," document JCTVC-I0408, JCTVC, Geneva, Switzerland, Apr. 2012.
- [2] F. Bossen, "Common HM test conditions and software reference configurations," document JCTVC-I1100, JCT-VC, Geneva, Switzerland, Apr. 2012.
- [3] A. Gabriellini, M. Mrak, D. Flynn, and M. Naccari, "Transform Skipping for Inter Predicted Coding Units," document JCTVC-J0077, JCT-VC, Stockholm, Sweden, Jul. 2012.

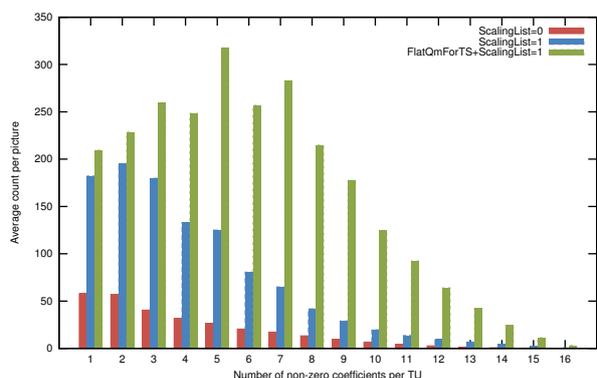
Figure 4: Effect of quantisation matrices on coefficient distributions of transform-skipped TUs (i_he10)



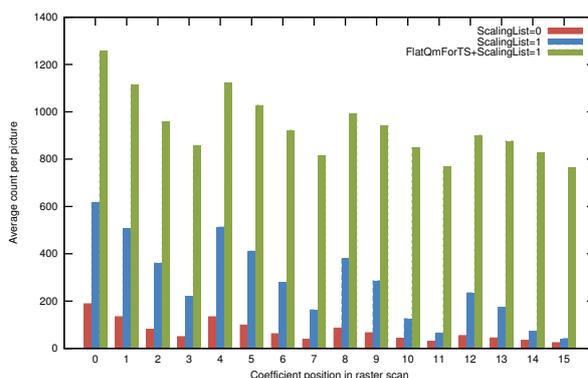
(a) Class A – Distribution of number of non-zero coefficients



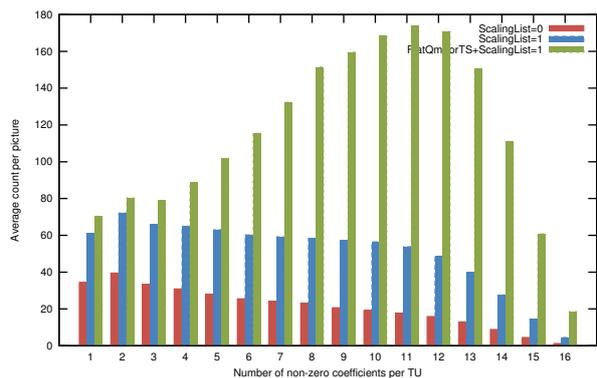
(b) Class A – Distribution of non-zero coefficients within a TU



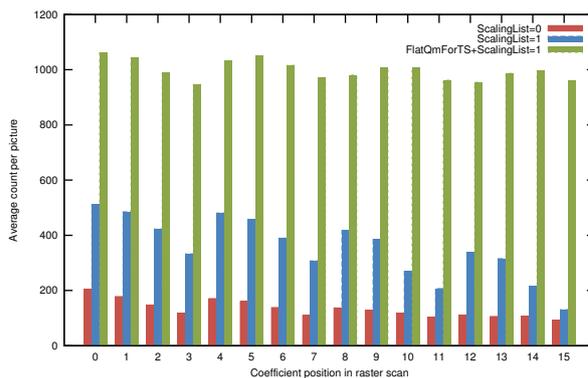
(c) Class B – Distribution of number of non-zero coefficients



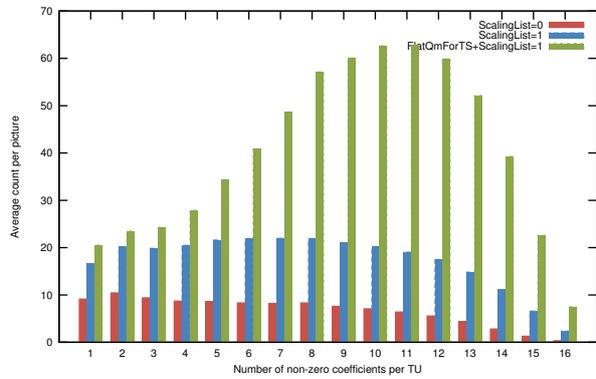
(d) Class B – Distribution of non-zero coefficients within a TU



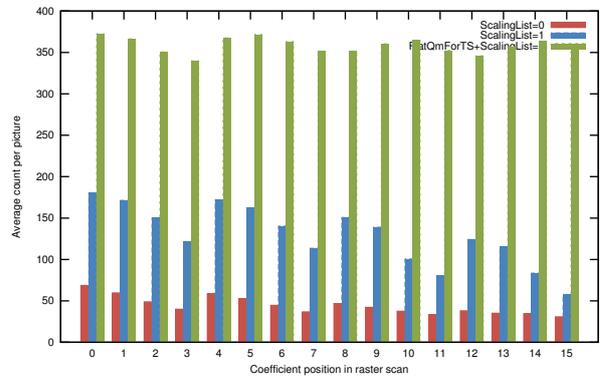
(e) Class C – Distribution of number of non-zero coefficients



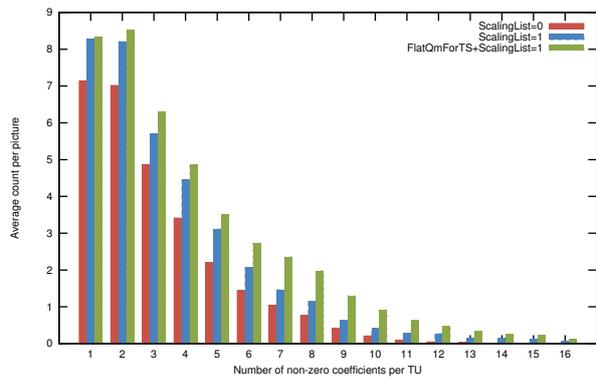
(f) Class C – Distribution of non-zero coefficients within a TU



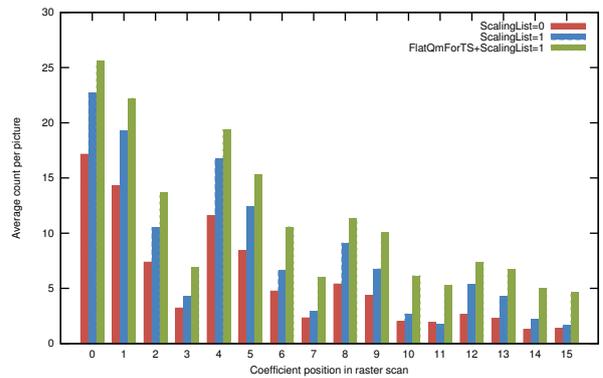
(g) Class D – Distribution of number of non-zero coefficients



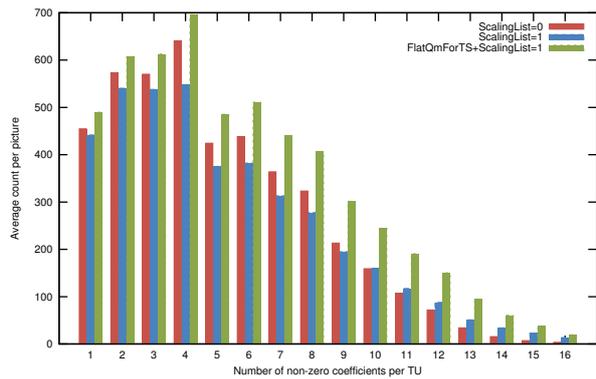
(h) Class D – Distribution of non-zero coefficients within a TU



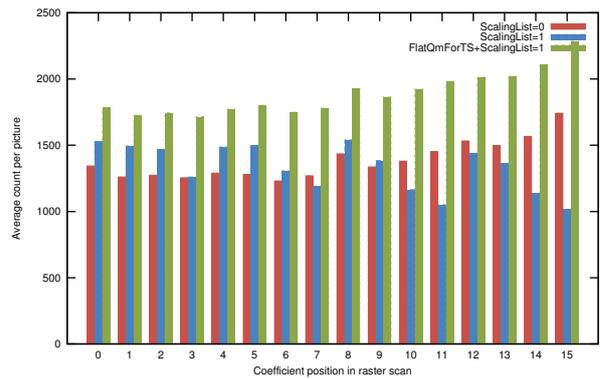
(i) Class E – Distribution of number of non-zero coefficients



(j) Class E – Distribution of non-zero coefficients within a TU



(k) Class F – Distribution of number of non-zero coefficients



(l) Class F – Distribution of non-zero coefficients within a TU