

INTERNATIONAL ORGANISATION FOR STANDARDISATION
ORGANISATION INTERNATIONALE DE NORMALISATION
ISO/IEC JTC1/SC29/WG11
CODING OF MOVING PICTURES AND AUDIO

ISO/IEC JTC1/SC29/WG11 MPEG/m52517
January 2020, Brussels, Belgium

<i>Source:</i>	Apple Inc.	
<i>Status:</i>	Input document	
<i>Title:</i>	CE13.29 report on in-loop geometry quantisation	
<i>Author(s):</i>	David Flynn Khaled Mammou	<i>davidflynn@apple.com</i> <i>kmammou@apple.com</i>

Abstract

This document examines the rate distortion behaviour of in-tree geometry quantisation as part of core experiment 13.29 [1, N18936].

Introduction

In-tree geometry scaling [2] provides a means to quantise (encoder) and scale (decoder) geometry positions while the coding tree is being constructed. The current software provides the following encoder modes:

- Whole slice scaling, where no per-node offsets are signalled, and the entire position value is scaled according to a per-slice QP. The QP is configurable.
- Per-node scaling, where the encoder chooses a per-node offset (a choice of three) based upon the node point density relative to other nodes in the level. The offsets are fixed at $\text{sliceQp} + \{-3, 0, 3\}$, however both sliceQp and the depth are configurable.

Analysis

In order to understand the effect of geometry position scaling, we determine the rate-distortion behaviour by varying QP at a given tree depth. To eliminate the effects of non-normative mode decisions, and to explore the fundamental behaviour, the software is modified as follows:

- The per-node QP selection based upon local density is removed so that all nodes have a QP offset of 0 and therefore a $\text{NodeQp} = \text{SliceQp}$.
- An encoder configuration parameter is added to select the node size at which quantisation is performed (rather than specifying the depth as per TMC13v8). This enables a consistent setting of the node size, which could otherwise vary between sequences and within slices of a sequence.

Evaluation is performed using two scaling node sizes: 2^5 and 2^7 . Simulations are performed over the range of QP from 5 to $4 + 6(\log_2(\text{nodeSize}) + 1) - 1$, inclusive. Larger QPs have no effect. Owing to the large number of test points, the following sequences are arbitrarily used: facade_00064_vox11, facade_00064_vox14, facade_00064_vox20, longdress_vox10_1300, loot_viewdep_vox12, overpass_q1mm, qnxadas-junction-approach, queen_0200, stanford_area_2_vox16, ulb_unicorn_vox20

Since quantisation is a lossy process, the lossy-geometry-lossy-attribute configuration is modified and used for this experiment, the CTC rate point specific configuration is removed and replaced as follows:

- Non-normative scaling related pre-processing (positionQuantizationScale) is disabled and replaced by the QP sweep.

- Attribute coding uses a fixed QP=34, corresponding to the CTC rate point r04.

A version of the software that implements the changes required for this test, along with configuration parameters is available in the [mpeg128/ce13.29/geomquant-nodeqpsweep](https://github.com/mpeg128/ce13.29/geomquant-nodeqpsweep) branch of the CE repository.

Bugfixes and common test condition deviations

It was discovered that the scaling function omits a clipping stage which, due to rounding in the scaling process, can produce a point outside its corresponding node. This was fixed in the CE branch and is described further in [3].

Due to interactions with both the planar mode and non-cubic geometry octrees, these features have been disabled in this experiment. The same related contribution suggests a means to resolve the interaction with QtBt.

Results

Figure 1 shows rate-distortion curves for four of the tested sequences. Further graphs are provided with this document.




Figure 1 – Rate-distortion behaviour of varying QP with scaling node sizes 2^5 (blue) and 2^7 (orange)

Observations

In some sequences there is a disparity between the rate-distortion performance of the 2^5 and 2^7 node sizes in sparse sequences (for instance fig. 1a). This appears to be due to IDCM nodes occurring prior to the quantisation depth and therefore being excluded from quantisation. A related contribution [4] proposes a method to address this effect.

In many instances the rate-distortion curve is characterised by severe steps. In many cases, the benefit of quantisation seems to be mostly related to removing a level from the octree, rather than increasing the density of the points. We continue to wonder if the use of the video-based QP mapping is appropriate. In extreme,

using step sizes of powers of two would greatly simplify implementation while retaining much of the power of the feature. Aspects of this are explored in a related contribution [5].

Furthermore, the effect on rate-distortion behaviour is not necessarily monotonic with respect to QP.

Finally, in-tree quantisation, while proposed for a particular type of content, may be used with any type of content. In particular, the feature may be used with dense voxelised sequences as found in category one. Figure 2 shows the visual effect of the feature when scaled with QP=11, resulting in a non-uniform spacing of the reconstructed points.




Figure 2 – Longdress with QP=11, rendered using an orthographic projection

References

- [1] 3DG, “CE 13.29 Geometry Quantization QP control,” ISO/IEC JTC1/SC29/WG11, 128th meeting, Geneva, Tech. Rep. w18936, Oct. 2019.
- [2] X. Zhang, W. Gao, S. Yea, and S. Liu, “[G-PCC][New proposal] Signaling delta QPs for adaptive geometry quantization in point cloud coding,” ISO/IEC JTC1/SC29/WG11, 127th meeting, Gothenburg, Tech. Rep. m49232, Jul. 2019.
- [3] D. Flynn and K. Mammou, “G-PCC: Minor simplifications and fixes to in-tree geometry quantisation,” ISO/IEC JTC1/SC29/WG11, 129th meeting, Brussels, Tech. Rep. m52521, Jan. 2020.
- [4] —, “G-PCC: Extension to in-tree geometry quantisation,” ISO/IEC JTC1/SC29/WG11, 129th meeting, Brussels, Tech. Rep. m52523, Jan. 2020.
- [5] —, “G-PCC: On in-tree geometry quantisation,” ISO/IEC JTC1/SC29/WG11, 129th meeting, Brussels, Tech. Rep. m52522, Jan. 2020.