

Point Cloud Reconstruction From Truncated Geometry-Based Streams

Diogo C. Garcia, *Senior Member, IEEE*, Andre L. Souto, Gustavo P. Sandri, *Member, IEEE*, Tomas M. Borges, Ricardo L. de Queiroz, *Fellow, IEEE*

Abstract—Geometry-based point cloud compression (G-PCC) has been rapidly evolving in the context of international standards. Despite the inherent scalability of octree-based geometry description, current G-PCC attribute compression techniques prevent full scalability for compressed point clouds. In this paper, we present a solution to add scalability to attributes compressed using the region-adaptive hierarchical transform (RAHT), enabling the reconstruction of the point cloud using only a portion of the original bitstream. Without the full geometry information, one cannot compute the weights in which the RAHT relies on to calculate its coefficients for further levels of detail. In order to overcome this problem, we propose a linear relationship approximation relating the downsampled point cloud to the truncated inverse RAHT coefficients at that same level. The linear relationship parameters are sent as side information. After truncating the bitstream at a point corresponding to a given octree level, we can, then, recreate the attributes at that level. Tests were carried out and results attest the good approximation quality of the proposed technique.

Index Terms—Point cloud compression, geometry-based point cloud compression, G-PCC, truncated bitstream.

I. INTRODUCTION

A Point cloud (PC) is a 3D structure usually represented by a collection of points or volume elements (voxels) described by their geometry, given as the (x, y, z) coordinates of the points, and by the points attributes, which may be color, normal vectors and reflectance, among others. Recently, PC compression (PCC) research has intensified [1]–[6] and the Motion Picture Expert Group (MPEG) is in the process of finalizing two PCC standards [7], [8], based on purely geometrical techniques (G-PCC) or on existing video compression standards (V-PCC). PCs can be used to represent 3D scenes, where the points describe the hull of objects therein, and are used in, for example, autonomous navigation [9], [10], heritage preservation [11], entertainment, and telepresence [12].

The wide range of applications may result in different requirements in terms of quality and resolution. Consider the example illustrated in Fig. 1, where a PC is combined with other 3D elements to compose a scene. If the PC is placed

near to the user’s viewing point, the composition of the scene would require it to be rendered at a higher resolution when compared to the PC being placed further away. On the user side, the PC can be downloaded and fully reconstructed and then downsampled to the required resolution. However, if the encoded bitstream could offer some degree of scalability, the user would be able to only download and process the required amount of the bitstream in order to reconstruct the PC at the desired resolution.

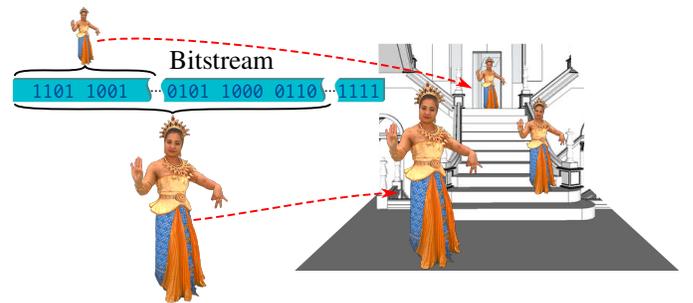


Figure 1: Reconstruction of the PC from a truncated bitstream. Subsequent portions of the bitstream are used to add details and resolution to the reconstructed PC.

Spatial scalability was one of the original requirements for G-PCC, so that the bitstream would have a layered structure for coarser approximations, and with each layer being used to predict the next one [13]. However, such a requirement is only partially addressed by G-PCC. The geometry representation already offers scalability by using octrees [14], wherein the point cloud geometry is inherently encoded as successive improvements on the geometry resolution, starting from a single block, successively dividing blocks into eight smaller blocks. The octree bitstream signals to the decoder which of these new blocks are occupied. For the attributes, G-PCC, which is currently based on RAHT [1], [2] or predicting/lifting transform [3], [15], offers partial spatial scalability, by using overlapping slices to encode regions of a PC at different fidelity levels, without inter-layer prediction [16].

In this paper, we propose a solution to reconstruct a PC from a truncated portion of the encoded bitstream containing the data from the octree-encoded geometry and RAHT-encoded attributes. The use of a truncated portion of the bitstream can save bandwidth in data transmission and avoid re-encoding the entire point cloud for each resolution.

Work partially supported by CNPq under grants 88887.600000/2021-00 and 301647/2018-6.

D. Garcia is with the Gama Engineering College, University of Brasilia, Brazil, e-mail: diogogarcia@unb.br

G. Sandri is with the Federal Institute of Brasilia, e-mail: gustavo.sandri@ifb.edu.br

R. de Queiroz is with the Department of Computer Science, University of Brasilia, Brazil, e-mail: queiroz@ieee.org

T. M. Borges and A. L. Souto are with the Electrical Engineering Department at University of Brasilia, Brazil, e-mails: tomas@divp.org and andre@image.unb.br.

II. TRUNCATED INVERSE RAHT

RAHT is a variation of the Haar transform, and it uses attribute values of a node at a lower level of the octree to predict the attributes of the nodes at the next level. For simplicity, assume there is just one attribute to be encoded per voxel (point), be it a color component, reflectance or else. Neighboring voxels are paired and transformed into low- and high-pass coefficients. The low-pass ones are further combined at each step with neighboring low-pass coefficients, repeating the process, until the entire space is traversed. At a given level, two low-pass coefficients about to be paired and transformed represent averages over different numbers of voxels, which render different weight values in the transformation matrix. Each weight indicates the number of voxels that were actually involved to generate that low-pass coefficient. Two neighbor low-pass coefficients at level $\ell + 1$, $F_{\ell+1,2n}$ and $F_{\ell+1,2n+1}$, are combined through an orthogonal transform to form a low- and a high-pass coefficient, $F_{\ell,n}$ and $G_{\ell,n}$, at level ℓ . Let $w_{\ell+1,2n}$ and $w_{\ell+1,2n+1}$ be the respective weights of the input coefficients, then

$$\begin{bmatrix} F_{\ell,n} \\ G_{\ell,n} \end{bmatrix} = \mathbf{T} \begin{bmatrix} F_{\ell+1,2n} \\ F_{\ell+1,2n+1} \end{bmatrix} ; \quad \mathbf{T} = \begin{bmatrix} a & b \\ -b & a \end{bmatrix}, \quad (1)$$

where

$$a = \sqrt{\frac{w_{\ell+1,2n}}{w_{\ell+1,2n} + w_{\ell+1,2n+1}}}, \quad b = \sqrt{\frac{w_{\ell+1,2n+1}}{w_{\ell+1,2n} + w_{\ell+1,2n+1}}}. \quad (2)$$

Note that $a^2 + b^2 = 1$, \mathbf{T} is orthogonal and equation (1) can be inverted using $\mathbf{T}^{-1} = \mathbf{T}^T$.

Let the PC have N occupied voxels, laid in a cubic grid of dimensions of 2^L voxels, which is referred to as an L -level PC. The set of $\{G_{\ell,n}\}$ are the $N - 1$ RAHT coefficients which are encoded along with the overall DC $F_{0,0}$. The forward RAHT start from the voxels (tree leaves at the L -th level, or $\{F_{L,n}\}$) generating low-pass coefficients which are laid in a voxel grid of level $L - 1$ ($\{F_{L-1,n}\}$). The process is recurred until we traverse all the way to the tree root at level 0, generating the overall DC value $F_{0,0}$ for the entire PC.

At the decoder, we start with $F_{0,0}$ and $G_{0,0}$, along with the weights $w_{1,0}$ and $w_{1,1}$, to calculate $F_{1,0}$ and $F_{1,1}$. From the start we need the weights for the whole tree. If we truncate the tree at level $L - K$ and know the geometry up to level $L - K$ we would still be unable to reconstruct the set of $\{F_{L-K,n}\}$ because the weights for all levels below $L - K$ still depend on the geometry at levels above $L - K$. Hence, RAHT is not scalable.

III. PROPOSED SOLUTION

Assume we have truncated data composed of the geometry information up to level $L - K$ and all the RAHT coefficients up to level $L - K$, i.e. $F_{0,0}$ and $\{G_{\ell,n}, 0 \leq \ell < L - K\}$. If we decode the data as an $L - K$ -level PC, the lower resolution geometry would provide incorrect weights $\{w'_{\ell,n}\}$ for the given coefficients, which were computed using the correct weights $\{w_{\ell,n}\}$. For example, all $w'_{L-K,n} = 1$ since it

is the last level of the truncated PC, which is most definitely not the case for the original PC with K further levels.

Let $X_{\ell,n}$ be the voxel attribute at the same position of $F_{\ell,n}$ obtained by downsampling the L -level PC from full-resolution voxels $\{F_{L,n}\}$ down to level ℓ . Generally, $X_{\ell,n}$ and $F_{\ell,n}$ are related by a scale depending on the weight $w_{\ell,n}$, and the set $\{X_{L-K,n}\}$ is what we want to reconstruct with the truncated bitstream. With the wrong weights w' , let the reconstructed voxel attributes at level $L - K$ be $\hat{F}_{L-K,n}$. We have noticed there is some correlation in between $X_{\ell,n}$ and $\hat{F}_{\ell,n}$ which is not precise but approximated. Figure 2 shows the relationship among pairs $X_{\ell,n}$ and $\hat{F}_{\ell,n}$ for different conditions for a couple of PCs encoded at particular bit-rates specified in MPEG's G-PCC common test conditions (CTC). Note the first order correlation among the variables. This pattern has held for all PCs and conditions we have tested in Sec. IV. This suggests a linear approximation from one variable to the other as

$$X_{\ell,n} \approx \alpha_{\ell} \hat{F}_{\ell,n} + \beta_{\ell}. \quad (3)$$

We, then, calculate, at the encoder side, not only the $\{F_{\ell,n}\}$ and $\{G_{\ell,n}\}$ for all levels $0 \leq \ell < L$, but we also calculate $\{\hat{F}_{\ell,n}\}$ and $\{X_{\ell,n}\}$ for $0 \leq \ell \leq L - K$. The encoder computes α_{ℓ} and β_{ℓ} using least-squares for all levels $0 \leq \ell \leq L - K$. The $2(L - K + 1)$ parameters are encoded and sent as side information to the decoder, which is generally a very small penalty.

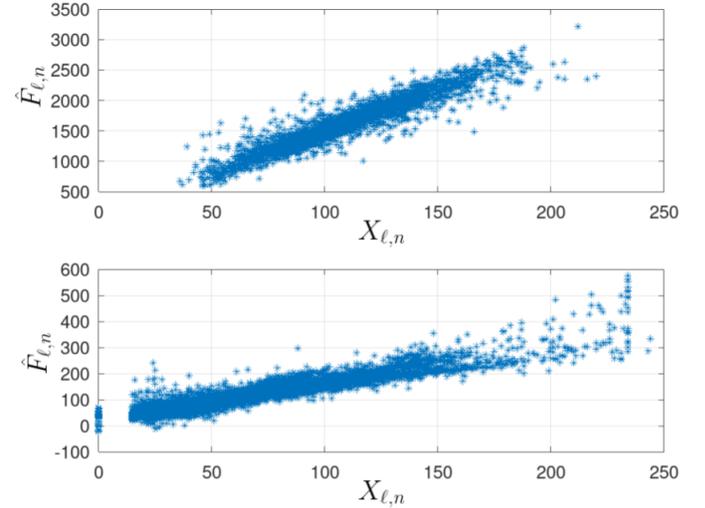


Figure 2: Relationship between attribute values $X_{\ell,n}$ of down-sampled point clouds and truncated inverse RAHT coefficients $\hat{F}_{\ell,n}$, for sequences *Longdress_vox_10*, frame 1300, CTC rate 1, downsampled by 16 (top plot), and *Ford_03_q1mm*, frame 0100, CTC rate 6, downsampled by 512 (bottom plot).

It is important to note that the G-PCC standard does not interlace the geometry and attribute information on a level basis, as each stream is separately included in the bitstream. In order to allow for the desired scalability in this standard, the proposed method expects stream interlacing on a level basis. The proposed G-PCC decoder must be able to truncate the geometry and attribute information until the desired

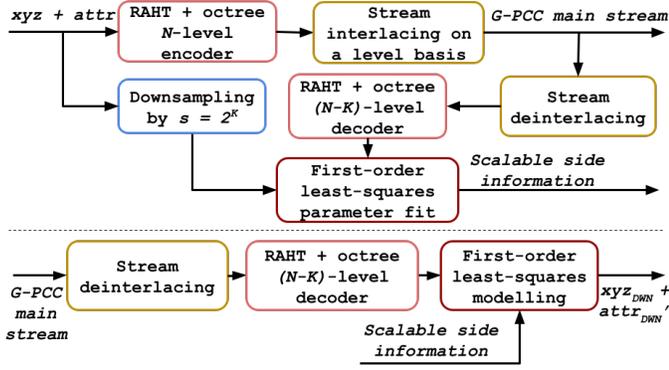


Figure 3: Truncated geometry-based encoder and decoder.

downsampling level, inverse transform the truncated RAHT coefficients, and then apply the first-order approximation. Figure 3 illustrates our method.

IV. EXPERIMENTAL RESULTS

In order to test the proposed solution, we selected PCs with different densities, levels, and attributes (RGB color or reflectance), as described in Table I. We tested different truncation points, corresponding to removing K octree levels from the point cloud, i.e., downsampling by a factor of $s = 2^K$. Our distortion metric computes Y-PSNR [17], [18], i.e. the PSNR of the luminance channel or reflectance. We start from a compressed bitstream and compare the resulting PC to the original PC after resolution reduction (downsampling) of K levels. The downsampling is carried out by removing octree levels of the geometry and by averaging the attributes at each level. Two methods are compared: the one obtained by full-stream reconstruction (*full-stream*) followed by down-sampling of K levels, against our algorithm to reconstruct $L - K$ levels of the PC from the truncated bitstream (*proposed*). The rate was calculated as the number of bits actually transmitted, i.e. the full bitstream size for *full-stream* method against the truncated bitstream size for the *proposed* one. Encoding and decoding, in both cases, were based on G-PCC Test Model (TMC13) version 12.0 [19], [20], and the encoding bit-rates were set according to the six target-rates from G-PCC’s CTC [21].

For a given PC, as we set K we can vary the encoding rate obtaining rate-distortion (RD) curves for each method (*full-stream* vs. *proposed*) from which one can calculate the Bjøntegaard-delta (BD) [22] rate reduction. Figure 4 shows the BD-rate achieved for many PCs at different truncation points.

Although impressive results were obtained, it is important to mention that reductions come with a caveat, since they do not capture the fact that a truncated bitstream cannot achieve the same distortion levels as full-stream decoding as we can see in the RD curves shown in Fig. 5. In yet another way we can present results, Fig. 6 relates the percentage of the bitstream that is used against the drop in Y-PSNR (dB) for the *proposed* method against the *full-stream* one. In these, one can see the curves for either varying the encoding rate for a given value of s (or K), or vice-versa. In all these curves one may appreciate there are sweet spots where the large reduction

Table I: Summary of information of the tested PCs.

Point cloud	L	# voxels	Attributes
(a) Longdress_vox_10_1051 [23]	10	765821	RGB
(b) Loot_vox_10_1000 [23]	10	784142	RGB
(c) Redandblack_vox_10_1450 [23]	10	729133	RGB
(d) Soldier_vox_10_0536 [23]	10	1059810	RGB
(e) ford_01_q1mm-0100	18	80265	Reflec.
(f) ford_02_q1mm-0100	18	79882	Reflec.
(g) ford_03_q1mm-0200	18	87590	Reflec.
(h) qnxadas-junction-approach_000001 [24]	18	31279	Reflec.
(i) qnxadas-junction-exit_000001 [24]	18	26842	Reflec.
(j) qnxadas-motorway-join_000001 [24]	18	29424	Reflec.
(k) qnxadas-navigating-bends_000001 [24]	18	26041	Reflec.

in rates and small reduction in quality may be of interest to given application developers.

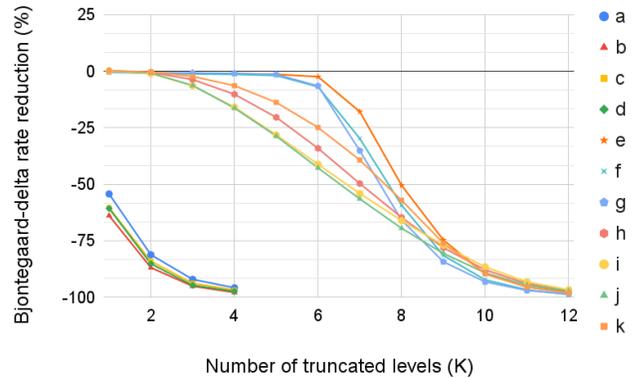


Figure 4: BD-rate reduction for several values of K (truncated levels) for several PCs, which are labeled as in Table I.

V. CONCLUSIONS

We have proposed an algorithm to reconstruct a G-PCC-coded point cloud from a partial (truncated) bitstream. Based on a linear relationship approximation relating the downsampled point cloud to the truncated inverse RAHT coefficients at that same level, the downsampled point cloud can be estimated, enabling lower-resolution “previews” without full-stream reconstruction. This is no substitute for a truly scalable coder, but it can be useful to save transmission bandwidth when, for example, rendering objects farther away. A number of tests were carried out to demonstrate the quality of the reconstruction. In many situations, one can save a large percentage of the original bit-rate at a small distortion penalty.

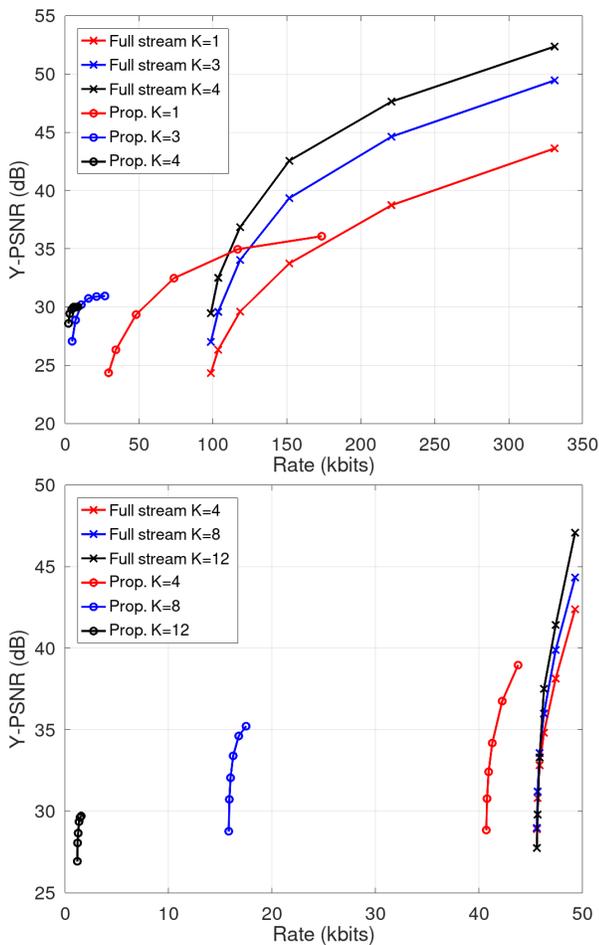


Figure 5: Representative RD curves for both methods at different levels of downsampling. Top plot is for *Longdress_vox10_1300* and the bottom one is for *qnxadas-junction-approach*.

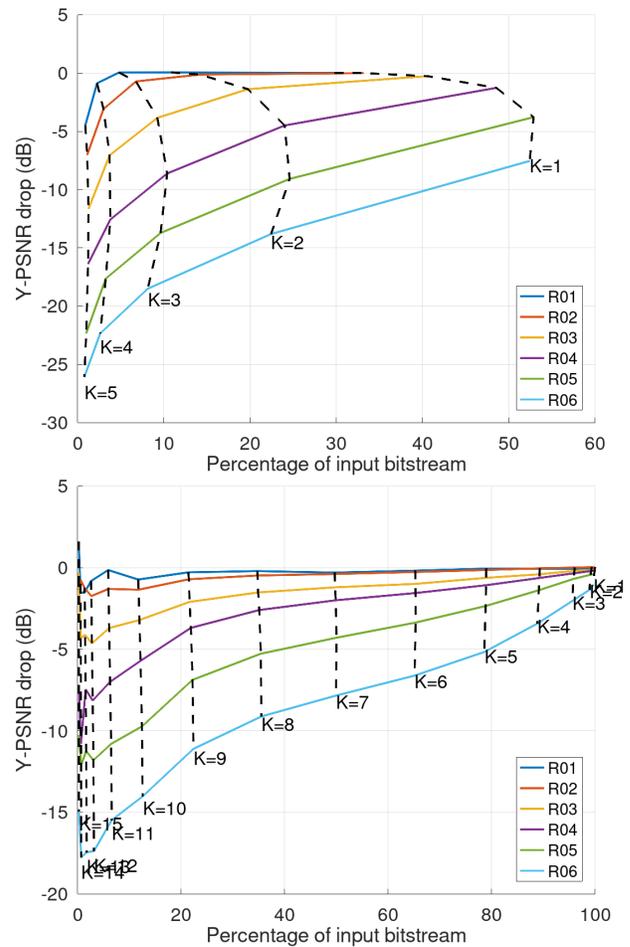


Figure 6: Bit-rate reduction and corresponding Y-PSNR drop for PCs *Longdress_vox10_1300* (top) and *qnxadas-junction-approach* (bottom), at different downsampling factors and encoding bit-rates. The dotted curves indicate different rates for the same K .

REFERENCES

- [1] R. L. de Queiroz and P. A. Chou, "Compression of 3D point clouds using a region-adaptive hierarchical transform," *IEEE Transactions on Image Processing*, vol. 25, no. 8, pp. 3947–3956, aug 2016.
- [2] S. Lasserre and D. Flynn, "On an improvement of RAHT to exploit attribute correlation," ISO/IEC MPEG JTC1/SC29/WG11, Geneva, CH, Tech. Rep. m47378, Mar. 2019.
- [3] K. Mammou, A. Tourapis, J. Kim, F. Robinet, V. Valentin, and Y. Su, "Proposal for improved lossy compression in TMC1," ISO/IEC MPEG JTC1/SC29/WG11, San Diego, US, Input contribution m42640, Apr. 2018.
- [4] G. Sandri, R. D. Queiroz, and P. A. Chou, "Compression of plenoptic point clouds using the region-adaptive hierarchical transform," in *2018 25th IEEE International Conference on Image Processing (ICIP)*. IEEE, oct 2018.
- [5] X. Zhang, P. A. Chou, M.-T. Sun, M. Tang, S. Wang, S. Ma, and W. Gao, "A framework for surface light field compression," in *2018 25th IEEE International Conference on Image Processing (ICIP)*. IEEE, oct 2018.
- [6] R. L. de Queiroz, C. Dorea, D. C. Garcia, R. U. Ferreira, D. R. Freitas, R. Higa, I. Seidel, and V. Testoni, "Differential plenoptic point cloud coding for V-PCC," ISO/IEC MPEG JTC1/SC29/WG11, Tech. Rep. input document WG7M55145, Oct. 2020.
- [7] D. Graziosi, O. Nakagami, S. Kuma, A. Zagherro, T. Suzuki, and A. Tabatabai, "An overview of ongoing point cloud compression standardization activities: video-based (V-PCC) and geometry-based (G-PCC)," *APSIPA Transactions on Signal and Information Processing*, vol. 9, p. e13, 2020.
- [8] S. Schwarz, M. Preda, V. Baroncini, M. Budagavi, P. Cesar, P. A. Chou, R. A. Cohen, M. Krivokuca, S. Lasserre, Z. Li, J. Llach, K. Mammou, R. Mekuria, O. Nakagami, E. Siahaan, A. Tabatabai, A. M. Tourapis, and V. Zakharchenko, "Emerging MPEG standards for point cloud compression," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 9, no. 1, pp. 133–148, 2019.
- [9] F. Hidalgo, "ORBSLAM2 and point cloud processing towards autonomous underwater robot navigation," in *Global Oceans 2020: Singapore – U.S. Gulf Coast*. IEEE, oct 2020.
- [10] M. Yoshioka, N. Sukanuma, K. Yoneda, and M. Aldibaja, "Real-time object classification for autonomous vehicle using LIDAR," in *2017 International Conference on Intelligent Informatics and Biomedical Sciences (ICIIBMS)*. IEEE, nov 2017.
- [11] D. Groux-Lecllet, J. Lentremy, G. Caron, and E. Mouaddib, "Efficient reproduction of heritage buildings: a new way to exploit 3d point cloud slicing : An example with the amiens gothic cathedral," in *2018 3rd Digital Heritage International Congress (DigitalHERITAGE) held jointly with 2018 24th International Conference on Virtual Systems & Multimedia (VSMM 2018)*. IEEE, oct 2018.
- [12] S. Orts-Escolano, C. Rhemann, S. Fanello, W. Chang, A. Kowdle, Y. Degtyarev, D. Kim, P. L. Davidson, S. Khamis, M. Dou, V. Tankovich, C. Loop, Q. Cai, P. A. Chou, S. Mennicken, J. Valentin, V. Pradeep, S. Wang, S. B. Kang, P. Kohli, Y. Lutchyn, C. Keskin, and S. Izadi, "Holoportation: Virtual 3d teleportation in real-time," in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, ser. UIST '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 741–754.
- [13] 3DG, "PCC Requirements," ISO/IEC MPEG JTC1/SC29/WG11, Gwangju, KR, Approved WG 11 document W17353, Jan. 2018.
- [14] D. Meagher, "Geometric modeling using octree encoding," *Computer Graphics and Image Processing*, vol. 19, no. 2, pp. 129–147, Jun 1982.
- [15] MPEG 3DG, "PCC Test Model Category 3 v0," ISO/IEC MPEG JTC1/SC29/WG11, Macau, CN, Approved N17249, Oct. 2017.
- [16] —, "Status on addressing the g-pcc requirements," ISO/IEC MPEG JTC1/SC29/WG11, virtual, MPEG Requirements N00052, Jan. 2021.
- [17] —, "MPEG 3DG and Requirements: Call for proposals for point cloud compression v2," ISO/IEC MPEG JTC1/SC29/WG11, Hobart, AU, Approved WG 11 document N16763, April 2017.
- [18] P. A. Chou and M. Krivokuca, "Concerns about Objective Metrics for Point Cloud Compression," ISO/IEC MPEG JTC 1/SC 29/WG 11, Macau, Input document to AhG on Point Cloud Compression m41809, October 2017.
- [19] 3DG, "G-PCC codec description v9," ISO/IEC MPEG JTC1/SC29/WG7, Online, Approved WG 7 document N00005, Nov. 2020.
- [20] —, "G-PCC test model v12 user manual," ISO/IEC MPEG JTC1/SC29/WG7, Online, Approved WG 7 document N00011, Oct. 2020.
- [21] —, "Common Test Conditions for G-PCC," ISO/IEC MPEG JTC1/SC29/WG11, Online, Approved WG 11 document N19584, July 2020.
- [22] G. Bjontegaard, "Calculation of average PSNR differences between RD curves," ITU, Tech. Rep. VCEG-M33, Apr. 2001.
- [23] E. d'Eon, B. Harrison, T. Myers, and P. A. Chou, "8i Voxelized Full Bodies, version 2 – A Voxelized Point Cloud Dataset," ISO/IEC JTC1/SC29 Joint WG11/WG1 (MPEG/JPEG), Geneva, input document m40059/M74006, January 2017.
- [24] D. Flynn, S. Lasserre, and G. Martin-Cocher, "PCC Cat3 test sequences from BlackBerry—QNX," ISO/IEC MPEG JTC1/SC29/WG11, Ljubljana, input document m23647, Jul. 2018.